

AD-A277 607



INTENTION PAGE	Approved for distribution	Form Approved public release 0704-0188
2		

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE		3. REPORT TYPE AND DATES COVERED FINAL/05 JAN 90 TO 04 JUL 93	
4. TITLE AND SUBTITLE APPLICATION OF GABOR REPRESENTATION TO MILITARY PROBLEMS (U)				5. FUNDING NUMBERS 7225/00 F49620-90-C-0016	
6. AUTHOR(S) Dr Richard Orr				7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Atlantic Aerospace Electronics Corp. 6404 Ivy Lane, Suite 300 Greenbelt, MD 20770-1406	
8. PERFORMING ORGANIZATION REPORT NUMBER AEOSR-TR- 94 0086				9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/NM 110 DUNCAN AVE, SUITE B115 BOLLING AFB DC 20332-0001	
10. SPONSORING / MONITORING AGENCY REPORT NUMBER F49620-90-C-0016				11. SUPPLEMENTARY NOTES 94-09430	
12a. DISTRIBUTION / AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE: DISTRIBUTION IS UNLIMITED Approved for public release distribution unlimited.				12b. DISTRIBUTION CODE UL	
13. SUMMARY In summary, the effort so far has proved in principle most of the supporting concepts, but has been insufficient to transition the work into the applications arena as yet. AAEC sees particular promise for this technology in certain applications areas, and is planning to propose effort in those areas. A key area is automatic target recognition (ATR). Machine -aided recognition problems have the feature that searching for objects can be enhanced in circumstances where shape characteristics of the objects are partially known in advance, either through a prior knowledge or data-aided algorithms. For example, in signal analysis, the Gabor transform is particularly adept at finding features having a common envelope. To maximally exploit such a circumstance, extraction of the analysis window from the data looms important. Given a large body of data such as that often encountered in an ATR problem, use of the data to drive the analysis functions seems wise as a measure to cut the amount of blind search, especially in view of findings that allegedly more "robust" tools such as the Wigner distribution can create artifacts through nonlinear processing if not used carefully. The role for optimum Gabor windowing in this scheme is clear, and as a result it appears that the best way in which to continue the line of work discussed above is to do it within the context of an application area such as ATR. The research at the point where it could profit from the interaction with real data as an aid in algorithm development/refinement. AAEC anticipates proposing a body of work of this nature as a logical follow-on to the work performed in both this contract and the cited SBIR's.					
14. SUBJECT TERMS				15. NUMBER OF PAGES 80	
16. PRICE CODE				17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	
18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED		19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED		20. LIMITATION OF ABSTRACT SAR(SAME AS REPORT)	

N 7540-01280-5500

Approved for public release;
distribution unlimited.

APPLICATION OF GABOR REPRESENTATION TO MILITARY PROBLEMS

Final Technical Report

2 July 1993

Sponsored by
Advanced Research Projects Agency (DoD)
ARPA Order No. 7225
Monitored by AFOSR Under
Contract # F49620-90-C-0016

Accession For	
NTIS	CRA&I <input checked="" type="checkbox"/>
DTIC	TAB <input type="checkbox"/>
Unannounced <input type="checkbox"/>	
Justification _____	
By _____	
Distribution /	
Availability Codes	
Dist	Avail and/or Special
A-1	

Atlantic Aerospace Electronics Corporation
6404 Ivy Lane, Suite 300
Greenbelt, MD 20770
Contract Date: 5 JAN 1990
Expiration Date: 4 JULY 1993
Reporting Period:
5 JAN 1992 - 4 JULY 1993

Principal Investigator:
Dr. Richard S. Orr
Phone: (301) 982-5215
Short Title of Work:
Final Technical Report

DTIC QUALITY INSPECTED 1

DISCLAIMER

"The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either express or implied, of the Defense Advanced Research Projects Agency or the U.S. Government."

TABLE OF CONTENTS

INTRODUCTION AND SUMMARY	1-1
1.1 INTRODUCTION	1-1
1.2 SUMMARY	1-1
1.2.1 Theoretical Developments	1-1
1.2.2 Software Development	1-3
1.2.3 Applications Research	1-3
THEORETICAL DEVELOPMENTS	2-1
2.1 THEORY	2-1
2.1.1 Gabor Sampling Theorem	2-1
2.1.2 Accuracy and Stability of the Expansion	2-5
2.1.3 Relations between Matrix Stability and Signal Duration	2-13
2.2 ALGORITHMS	2-15
2.2.1 Oversampling with the Gabor Transform	2-15
SOFTWARE DEVELOPMENT	3-1
3.1 SUMMARY OF SECOND YEAR EFFORT	3-1
3.1.1 Phase II Code	3-1
3.1.2 Capabilities	3-3
3.2 THIRD YEAR EFFORT	3-5
3.2.1 New Facilities	3-5
3.2.2 Phase III Code	3-5
3.2.3 Warnings and Suggestions	3-7
3.2.4 Scripts	3-7
3.3 USERS GUIDE	3-8
3.3.1 Procedure Descriptions	3-8
3.3.2 Program Flow	3-9
APPLICATIONS RESEARCH	4-1
4.1 FAULT IDENTIFICATION IN FEEDBACK CONTROL CIRCUITS	4-1
4.1.1 Summary of Previous Work (Greenbelt Facility)	4-1
4.1.2 Summary of Previous Work (Waltham Facility)	4-2
4.1.3 Year 3 Effort (Waltham Facility)	4-3
4.1.3 Year 3 Effort (Greenbelt Facility)	4-21
4.1.5 Future Work	4-22
4.2 AUTOMATIC TARGET RECOGNITION	4-23
4.2.1 Problem Definition	4-23
4.2.2 Experimental Results	4-27
4.2.3 Future Work	4-49
4.3 MINIMUM DIMENSION GABOR	4-50
4.3.1 Summary of SBIR Contract Effort	4-50
4.3.2 Extensions Under Current Contract	4-55
4.3.3 Future Work	4-72
REFERENCES	A-1
STATEMENT OF WORK	B-1
PERIODIZED GABOR PAPER	C-1
MINIMUM DIMENSION GABOR PAPER	D-1

FIGURES

Fig. 2.1-1. Graph of the periodic variance of σ .	2-12
Fig. 3.1-1. Logical flow of the GPS package at the end of the second year	2-4
Fig. 3.3-1. Illustration of program flow.	3-10
Fig. 4.1-1. Typical raw signals (plotted with amplitude offsets).	4-5
Fig. 4.1-2. Classifier system block diagram.	4-5
Fig. 4.1-3. Erosion process.	4-6
Fig. 4.1-4. Signal conditioning experiments.	4-6
Fig. 4.1-5a. Effect of signal conditioning	4-7
Fig. 4.1-5b. Effect of signal conditioning	4-8
Fig. 4.1-5c. Effect of signal conditioning	4-8
Fig. 4.1-6. Calculation of pattern spectrum.	4-9
Fig. 4.1-7. Feature extraction experiments.	4-10
Fig. 4.1-8. Example of log-spaced pattern spectra.	4-11
Fig. 4.1-9. Maximum-likelihood classification.	4-11
Fig. 4.1-10. Emitter classification experiments.	4-12
Fig. 4.1-11. Sole misclassified signal (mmm3mmme). (a) Versus typical signal from true class. (b) Versus typical signal from estimated class	4-14
Fig. 4.1-12. Classifier performance evaluation using class reassignment.	4-18
Fig. 4.2-1. Cloud cover section.	4-27
Fig. 4.2-2. Targets.	4-28
Fig. 4.2-3. Clouds plus targets.	4-28
Fig. 4.2-4. Unfiltered signal.	4-29
Fig. 4.2-5. Filtered signal containing the targets.	4-29
Fig. 4.2-6. Longwave response used as a window.	4-30
Fig. 4.2-7. Gabor coefficients at 0 shifts.	4-30
Fig. 4.2-8. Gabor coefficients at 1 shifts	4-31

Fig. 4.2-9. Gabor coefficients at 2 shifts.....	4-31
Fig. 4.2-10. Gabor coefficients at 3 shifts.....	4-32
Fig. 4.2-11. Processed Gabor coefficients.	4-33
Fig. 4.2-12. Signal reconstructed from the nonlinearly processed coefficients.	4-34
Fig. 4.2-13. Cloud unfiltered.....	4-35
Fig. 4.2-14. Cloud filtered.....	4-35
Fig. 4.2-15. Gabor coefficients at 0 shifts.....	4-36
Fig. 4.2-16. Gabor coefficients at 1 shift.	4-36
Fig. 4.2-17. Gabor coefficients at 2 shifts.....	4-37
Fig. 4.2-18. Gabor coefficients at 3 shifts.....	4-37
Fig. 4.2-19. Clouds with non-uniformly spaced targets.....	4-38
Fig. 4.2-20. Gabor coefficients at 0 shifts.....	4-38
Fig. 4.2-21. Gabor coefficients at 1 shift.	4-39
Fig. 4.2-22. Gabor coefficients at 2 shifts.....	4-39
Fig. 4.2-23. Gabor coefficients at 3 shifts.....	4-40
Fig. 4.2-24. Mediumwave response used as a window.....	4-40
Fig. 4.2-25. Gabor coefficients at 0 shifts: MW window.	4-41
Fig. 4.2-26. Gabor coefficients at 1 shift: MW window.	4-41
Fig. 4.2-27. Gabor coefficients at 2 shifts: MW window.	4-42
Fig. 4.2-28. Gabor coefficients at 3 shifts: MW window.	4-42
Fig. 4.2-29. Signal resulting from the removal of the negative values of the original filtered signal.	4-43
Fig. 4.2-30. Gabor coefficients for the processed signal with $M = 512$ and $N = 2$. (a) Coefficients for the unshifted signal and (b) coefficients for the signal after time shifting once to the left.....	4-44
Fig. 4.2-31. Gabor coefficients for the processed signal with $M = 256$ and $N = 4$. (a) - (d) correspond to the coefficients for the signal after time shifting 0 - 3 points to the left respectively.	4-46
Fig. 4.2-32. Oversampled Gabor coefficients for the processed signal with $M = 1024$ and $N = 2$. (a) Coefficients for the unshifted signal and (b) coefficients for the signal after time shifting once to the left.....	4-47

Fig. 4.2-33. Oversampled Gabor coefficients for the processed unshifted signal with $M = 1024$ and $N = 4$.	4-48
Fig. 4.2-34. Oversampled Gabor coefficients for the processed unshifted signal with $M = 1024$ and $N = 8$.	4-48
Fig. 4.3-1 Nonlinear constraint that controls exponential growth of the biorthogonal on the basis of the largest magnitude value within each Gabor time slice.	4-57
Fig. 4.3-2 Nonlinear constraint that controls exponential growth of the biorthogonal on the basis of the magnitudes point-by-point within each Gabor time slice.	4-58
Fig. 4.3-3. Nonlinear constraint that controls range within the first Gabor time interval.	4-59
Fig. 4.3-4. Linear constraint that controls exponential growth of the biorthogonal on the basis of the largest magnitude value within the first and last Gabor time slice.	4-60
Fig. 4.3-5. Transient signal - bounds, no constraints	4-61
Fig. 4.3-6. Transient signal - bounds, and nonlinear constraints	4-62
Fig. 4.3-7. Initial biorthogonal and corresponding window.	4-64
Fig. 4.3-8. Two signals - no optimization, window equals rectangular pulse.	4-65
Fig. 4.3-9. Two signals - no optimization, window equals exponential pulse.	4-66
Fig. 4.3-10. Two signals- bounds, no constraints.	4-67
Fig. 4.3-11. Two signal case - bounds, linear constraints.	4-68
Fig. 4.3-12. Two signal case - bounds, nonlinear constraints.	4-69
Fig. 4.3-13. Two signal case - bounds, nonlinear constraints.	4-70
Fig. 4.3-14. Two signal case - bounds, nonlinear constraints.	4-71

TABLES

Table 4.1-1 - MINDIS classifier results for old data set. (a) Using mean or median feature vectors for each class. (b) Using first feature vector for each class.	4-13
Table 4.1-2. MINDIS classifier results for extended old data set.	4-15
Table 4.1-3. Tallies of sponsor-defined class members in each neural-net-defined class.	4-19
Table 4.1-4. Distribution of population over neural net classes.	4-20
Table 4.1-5. MINDIS classifier performance relative to neural net class assignments.	4-20
Table 4.1-6. Exploration of alternate feature extraction methods.	4-21

CHAPTER 1

INTRODUCTION AND SUMMARY

1.1 INTRODUCTION

This document constitutes the Final Technical Report for contract F49620-90-C-0016 Application of Gabor Representations to Military Problems. All work performed by Atlantic Aerospace Electronics Corporation under the Final year of this contract is reported herein.

The outline of this report is as follows. The remainder of this chapter contains short summaries of the work items. Chapter 2 reviews theoretical developments, Chapter 3 discusses the software tasks and Chapter details the applications areas investigated. There are three of these: (4.1) application of morphology to specific emitter analysis; (4.2) use of Gabor transforms in automatic target recognition and (4.3) minimum dimension Gabor representations.

1.2 SUMMARY

1.2.1 Theoretical Developments

Discussion of theory is presented in two parts: analytical results; and algorithms.

1.2.1.1 Theory

Theory topics investigated include a Gabor sampling theorem, further studies on accuracy and stability of the representation and relations between matrix method stability and signal duration.

1.2.1.1.1 Gabor Sampling Theorem

In section 2.1.1 we show a sampling theorem applicable to Gabor expansions of bandlimited functions, derived rather directly by Shannon's technique. The theorem illuminates relationships that exist between the number of degrees of freedom in the signal and window functions and the minimum number of Gabor coefficients that must be retained to have a faithful expansion.

1.2.1.1.2 Accuracy and Stability

In section 2.1.2 we report on the stabilization of the Gabor transform under situations where a window choice desirable from an application point of view leads to a near singular mapping.

Methods applicable to both versions of the transform used in GSPS (Zak and matrix) are presented. In addition, some thought has been given to the use of discrete representations in which the window function has a zero in its Zak transform, this being one of the singular situation for which the techniques mentioned above were devised. The philosophy of such methods is to live with the zero rather than modify it. One technique is to represent the signal as a projection into the space of all signals having Zak zeroes where the window function has them, and minimizing the energy loss by through relative positioning of the signal and grid. The error involved in neglecting—in Zak space—a coordinate for which the Zak transform of the window is very small or zero is considered, and the minimization of this error over shifts of the signal and window is shown to lead to smaller than errors than one might initially expect. A specific example—white gaussian noise input—is carried through the analysis. Further improvement is possible when the input is cyclostationary.

We also indicate how a slightly modified basis, one that eliminates a number of the Gabor translates from the basis—as many as there are Zak zeroes in the window—and replaces them with other functions, provides a stable basis that retains most of the Gabor features.

1.2.1.1.3 Relations between Matrix Stability and Signal Duration

In section 2.1.3, stability of the matrix method of Gabor transform calculation is studied by looking at the algorithm that produces a biorthogonal function from its window, or vice versa. Our preliminary results help explain the mechanisms controlling the blowup of the transform method for certain window selections.

1.2.1.2 Algorithms

1.2.1.2.1 Oversampling with the Gabor Transform

One new important algorithm has been implemented in GSPS this year, which is the oversampled Gabor representation, or the Weyl-Heisenberg expansion. In this theory the number of expansion coefficients exceeds the number of data samples, which leads to some interesting tradeoffs. For the higher computation required to obtain the coefficients, one gets finer time-frequency resolution—as in classical Gabor, one can partition this between domains at will. As could be expected, this induces a non uniqueness in the transform, which one can view in one of two equivalent ways: (1) the coefficients associated with a fixed function are not unique, or (2) the relationship between window and biorthogonal is no longer uniquely invertible.

1.2.2 Software Development

1.2.2.1 Second Year Effort

The accomplishments of the second year effort, during which the bulk of the GSPS was developed, are reviewed prior to discussing the third year amendments.

1.2.2.2 Third Year Effort

Among the principal upgrades of the GSPS during year three are: a map view of the coefficients; ability to turn grid lines off or on at will; improvement of the command line interface; and availability of intermediate results in the transforms. In accomplishing the latter, the transforms were split so that, for example, the computation can be halted after the calculation of the Zak transform without continuing through the Gabor coefficient calculation. Also, a list of warnings is provided about features of the code for which undebugged traps have been observed.

1.2.2.3 User's Guide

New files and routines are listed, and as an example, the procedure for adding a new signal type is explained. Because of the split-up of the transform calculations, a new program flow is in operation, and this is discussed in detail.

1.2.3 Applications Research

1.2.3.1 Fault Identification in Feedback Control Circuitry

In work performed during the first two years of the contract, AAEC investigated the ability to locate faults in feedback control circuitry using data sets provided by a DOD customer. In this last year the emphasis was on the test of morphological filter methods for this problem, under funds added to the contract through a MIPR. A significant issue in the evaluation of the methods was the questionable "ground truth" supplied with the data. Near the end of the work a revision of the ground truth in the data was made using a neural net classifier, and subsequent evaluations were made against this standard. Classification using the methods described in the text generally succeeded in achieving a high correct classification probability.

An extension of the Gabor methods used on this problem in year 2 was attempted, but the "ground truth" problems associated with the data precluded any meaningful effort, since the revised ground truth was not available in time for the needed analysis.

1.2.3.2 Automatic Target Recognition

An infrared search and track (IRST) system is able to detect targets while they are still unresolved by the receive optics. Each sufficiently strong target generates a response equal to the point spread function (spatial impulse response) of the optics. The premise that automatic detection and recognition (ATR) of such targets in a background of, for example, clouds could be enhanced by Gabor-based signal processing was tested under contract.

AAEC employed some actual cloud background data upon which was superimposed both uniformly and nonuniformly distributed targets representing point spread functions of either a long wave or medium wave IR sensor. Single scan lines were extracted and subjected to Gabor processing characterized by: (1) a window function modeled after the target impulse response, and (2) expansion that emphasized time (or spatial) resolution at the expense of frequency resolution.

The experimental results showed two main features. The first is that if one performs not a single Gabor analysis, but instead several in which the registration of signal and window are changed at each instance, the targets essentially all show up in one of the cuts. In some other cases in which preprocessing was applied to the signals, the targets were made visible in a single cut, but spread across the frequency bins.

The results of this investigation are of course preliminary and incomplete, but they do suggest promise for the Gabor transform as a tool in this class of ATR problems. Applications involving resolved targets were not tested.

1.2.3.3 Minimum Dimension Gabor

Under a prior Phase I and II SBIR for ARPA, AAEC has performed the theoretical and initial computational exercises in minimization of the dimension of a representation of a signal set. The efforts carried out under this contact further generalize this by incorporating minimization under linear and nonlinear constraints on the biorthogonal function. The results of the exercise show this process to be tricky, requiring greater effort than was able to be applied here. The computational procedures in doing this have by now been pretty well wrung out, but there remain fundamental questions associated with the best types of constraints to use. Also, the ability to work with the Zak method, which we believe to be the more stable of those available, has yet to be implemented.

CHAPTER 2

THEORETICAL DEVELOPMENTS

2.1 THEORY

2.1.1 Gabor Sampling Theorem

During the reporting period we have looked at the application of sampling theorems within the realm of Gabor analysis. This activity has led to a Gabor sampling theorem that casts light on the relationship between degrees of freedom in the Gabor window and the represented signal. We begin with a naive imitation of the original Shannon-Kotelnikov theorem derivation and encounter some interesting results that differ from the classical theory due to the supervening circumstance that there are two functions of the independent variable (taken to be time here) subjected to sampling. This imposes some stipulations not present in the earlier theory.

To begin, we address the use of Shannon-Kotelnikov sampling expansions for functions expressed in a Gabor representation. In general, a function represented as

$$s(t) = \sum_{m,n} a_{m,n} w_{m,n}(t), \quad (2.1.1)$$

where the $\{w_{m,n}(t)\}$ comprise a Gabor basis for some $T > 0$,

$$w_{m,n}(t) = w(t - nT) \exp(j2\pi mt / T), \quad (2.1.2)$$

does not admit of a sampling theorem because it is not necessarily bandlimited. Even if the window function w is itself bandlimited, (2.1.1) can have arbitrarily high frequency content if the Gabor coefficients do not vanish as m increases. But with a bandlimited window and Gabor coefficients that are zero beyond a maximum frequency, any linear combination of basis functions will be bandlimited, hence subject to a sampling theorem. We develop such a theorem and use it, along with assumptions of time limitation, to investigate relationships among the time-bandwidth products of the signal and window and the number of nonzero time and frequency coefficients in the expansion.

2.1.1.1. The Sampling Theorem

Let $w(t)$ be a function bandlimited to $|f| < B_w$, where f is cyclic frequency, such that the set of translates (2.1.2) is a basis for $L^2(\mathbb{R})$. We denote the set of such bandlimited functions as $B\{B_w\}$ and say that $w \in B\{B_w\}$. Correspondingly, the translate $w_{m,n} \in B\{B_w + m/T\} \forall n \in \mathbb{Z}$. For an expression of the form

$$s(t) = \sum_{m=-M/2}^{M/2} \sum_{n=-\infty}^{\infty} a_{m,n} w_{m,n}(t) \quad (2.1.3)$$

we clearly have that $\{w_{m,n}\} \in B\{B_w + M/2T\} \forall n \in \mathbb{Z}, m \leq M/2$, and, by Shannon-Kotelnikov [SHAN],

$$\begin{aligned} w_{m,n}(t) &= \sum_{k=-\infty}^{\infty} w_{m,n}\left(\frac{k}{2B_0}\right) \text{sinc}\left[2\pi B_0\left(t - \frac{k}{2B_0}\right)\right] \\ &= \sum_{k=-\infty}^{\infty} w\left(\frac{k}{2B_0} - nT\right) \exp(j2\pi nk / 2B_0 T) \text{sinc}\left[2\pi B_0\left(t - \frac{k}{2B_0}\right)\right], \end{aligned} \quad (2.1.4)$$

where $B_0 = B_w + M/2T$ and we use $\text{sinc}(\cdot)$ to denote $\sin(\cdot)/(\cdot)$. To use a common set of samples of w for each value of n in (2.1.4), we should set T equal to a multiple of $1/2B_0$, i.e., $T = L_0 / 2B_0$. Now (2.1.4) reads

$$w_{m,n}(t) = \sum_{k=-\infty}^{\infty} w\left(\frac{k - nL_0}{2B_0}\right) \exp(j2\pi nk / L_0) \text{sinc}\left[2\pi B_0\left(t - \frac{k}{2B_0}\right)\right]. \quad (2.1.5)$$

Inserting the sampling theorem expansion for w into (2.1.2) yields

$$s(t) = \sum_{m=-M/2}^{M/2} \sum_{n=-\infty}^{\infty} a_{m,n} \sum_{k=-\infty}^{\infty} w\left(\frac{k - nL_0}{2B_0}\right) \exp(j2\pi nk / L_0) \text{sinc}\left[2\pi B_0\left(t - \frac{k}{2B_0}\right)\right]. \quad (2.1.6)$$

Trivially $s \in B\{B_0\}$, and it has the sampling expansion

$$s(t) = \sum_{k=-\infty}^{\infty} s\left(\frac{k}{2B_0}\right) \text{sinc}\left[2\pi B_0\left(t - \frac{k}{2B_0}\right)\right]. \quad (2.1.7)$$

where

$$s\left(\frac{k}{2B_0}\right) = \sum_{m=-M/2}^{M/2} \sum_{n=-\infty}^{\infty} a_{m,n} w\left(\frac{k-nL_0}{2B_0}\right) \exp(j2\pi nk / L_0). \quad (2.1.8)$$

2.1.1.2 Time-Bandwidth Products

No restriction on the time index n of the Gabor coefficients was needed to get the sampling theorem. But now let us apply the restriction $|n| < N/2$ and consider only those expansion using a finite number of coefficients:

$$s(t) = \sum_{m=-M/2}^{M/2} \sum_{n=-N/2}^{N/2} a_{m,n} w_{m,n}(t). \quad (2.1.9)$$

Further suppose that w is essentially time-limited, i.e., the samples of w are zero for $|t| \geq T_w / 2$. We can associate with w a time-bandwidth product (or number of degrees of freedom)

$$D_w = 2B_w T_w, \quad (2.1.10)$$

that represents the minimum number of samples of required to represent w faithfully through the sampling theorem. The same can be done for s , and it is easy to show that

$$D_s = 2B_s T_s = (T_w + NT)(2B_w + M/T). \quad (2.1.11)$$

Because both $L_0 (= 2B_0 T)$ and M are integers, and $L_0 = 2B_w T + M$, the quantity $2B_w T$ is also an integer, which we denote as L , the length of the Gabor time cell in units of Nyquist samples of s . Then

$$D_s = (1 + M/L)(2B_w T_w + NL). \quad (2.1.12)$$

Assume the window function and the coefficient bounds M and N are given, and that only the grid shape (L) may be varied. To varying L is to change the subspace of $L^2(\mathbb{R})$ spanned by the basis of the truncated representation. Although we would like our basis to accommodate as many degrees of freedom as possible, it is not particularly desirable to have that number depend

on the grid shape. A value of time-bandwidth product that is achievable for any grid choice would be a conservative guide to assessing the capability of the truncated representation.

Minimizing the value of D_s with respect to L , however, determines a value of signal time-bandwidth product that is achievable independent of the choice of grid. Signals whose time-bandwidth products are less than this minimum will be called uniformly representable with respect to the window and the coefficient limits $\{w, M, N\}$. D_s has a unique minimum at $L = L^*$,

$$D_s|_{L=L^*} = [\sqrt{2B_w T_w} + \sqrt{MN}]^2 \quad (2.1.13)$$

$$L^* = \sqrt{\frac{M}{N} 2B_w T_w}. \quad (2.1.14)$$

when $L^* \geq 1$. We interpret (2.1.14) to say that a signal is uniformly representable by the basis $\{w, M, N\}$ if its time-bandwidth product satisfies

$$\sqrt{2B_s T_s} \leq \sqrt{2B_w T_w} + \sqrt{MN}. \quad (2.1.15)$$

If (2.1.15) does not hold for some s , then either the time extent of the signal will exceed the span of the time translates of w or else s will have frequency content that is not captured by the frequency translates, at least for some choices of T (or L). One of the following inequalities will be violated for some choices of grid:

$$\begin{aligned} B_s &> B_w + M/2T \\ T_s &> T_w + NT \end{aligned} \quad (2.1.16)$$

Although no theorem guarantees that a signal uniformly representable with respect to a truncated basis $\{w, M, N\}$ will automatically have a 'nice' representation in the sense of some error measure, for example

$$\varepsilon = \left\| s(t) - \sum_{m,n} a_{m,n} w_{m,n}(t) \right\|_2, \quad (2.1.17)$$

the time-bandwidth product guidelines should provide insight into the match between a signal and its expansion over a truncated basis. Theorems relating to this will be sought in future work.

We close by observing that (2.1.15) can be solved in the following two interesting forms: the first states a requirement on the window time-bandwidth product, given the signal and the number of coefficients; the second tells how many coefficients must be retained to expand a given function with a particular window:

$$\begin{aligned}\sqrt{2B_w T_w} &\geq \sqrt{2B_s T_s} - \sqrt{MN} \\ \sqrt{MN} &\geq \sqrt{2B_s T_s} - \sqrt{2B_w T_w}.\end{aligned}\tag{2.1.18}$$

2.1.2 Accuracy and Stability of the Expansion

Accuracy and stability of the Gabor representation are longstanding issues. Over the course of work on the three-year contract, AAEC has made considerable progress in understanding of these matters, most of which has been discovered piece by piece and not documented in prior reports.

Two distinct approaches are reported here. In the first type, a "bad" window is tamed by adjustments to either the window or the point grid. Finding a transformation close to the original, but with much better conditioning, is the goal of this approach. In a second method, we examine the implications of working with window functions that lead to noninvertible transforms and assessing the associated loss of representation fidelity. The latter work is in its early stages and is not yet supported by numerical experiments, whereas the former category has been well explored numerically by AAEC.

2.1.2.1 Stabilization by Window or Grid Adjustment

In [BALA], presented at the 1992 IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis, we have collected descriptions of a number of the techniques we have successfully employed in stabilizing ill-conditioned Gabor expansions and presented them by computed example. Methods applying to both the Zak and matrix methods are included. The following summarizes the content of this paper, which is found in full in Appendix B.

The paper begins by summarizing the Zak and matrix algorithms used in AAEC's GSPS software, highlighting the features of the window function that in each case can lead to singularity of the transform. Section 2 addresses experiments performed with the Zak algorithm, where we present a condition number expression for the mapping. Explosion of this number

results from zeros of the Zak transform of the window function lying on or near the computation grid. Using a window of the form

$$w(t) = te^{-at}u(t), \quad (2.1.19)$$

which has a single zero at a point whose time coordinate is adjustable by choice of a , we explore three techniques for ameliorating the Zak zero problem: (1) subtle change of the window waveform to displace the Zak zero; (2) time- or frequency-translation of the window; and (3) alteration of the time-frequency grid. Examples of an ill-conditioned transform before and after each of these methods are shown, and in each case considerable stability is recovered.

Section 3 deals with a similar treatment of the matrix method. The failure mechanism in this case is the occurrence of zeros of the window function itself within the interval from the zero-th to the first time grid point. The window in (2.1.19), having a zero at the time origin, makes its matrix transform noninvertible. Left shifting the window by just one point is enough to stabilize the transform. A second method involves adding a small constant value to the window at every point. Numerical results of both techniques are presented.

2.1.2.2 Working with Transforms Having Zak Zeros

The following thoughts were stimulated by reading [TOLJ], in which it is remarked that although one approach to deal with a window having a Zak zero is to use it only for expansion of functions that have a corresponding Zak zero, this solution may be overly restrictive. The reference contains no evaluation of the degree of restriction imposed by this approach, and it seemed of interest to consider this aspect further.

Assume a discrete Gabor expansion using a window having exactly one Zak zero on the grid, e.g., the gaussian, for which

$$Z_w(\frac{1}{2}, \frac{1}{2}) = 0 \quad (2.1.20)$$

In a discrete time expansion there is just one bad point for the transform as executed by the Zak method, which is the point where the "divide by zero" occurs. Consider the set of all functions f in ℓ^2 such that

$$Z_f(\frac{1}{2}, \frac{1}{2}) = 0. \quad (2.1.21)$$

This set is a subspace of ℓ^2 , since if f and g have the zero property, $af + bg$ has it also. We would like to consider what it means to take an arbitrary f and make some "minimal" modification to it such that the resulting \hat{f} belongs to the subspace $Z_{(f)}(\frac{1}{2}, \frac{1}{2}) = 0$.

The indicated procedure is of course a projection operator, and thus there is a least squares solution to the estimation. That is, find \hat{f} such that: (i) $Z_{(\hat{f})}(\frac{1}{2}, \frac{1}{2}) = 0$, and (ii) $\|\hat{f} - f\|^2$ is minimized. Finding \hat{f} could be accomplished by a classical least squares procedure, but there is a better way. One should compute the Zak transform of f , zero the appropriate coordinate, and retransform to get \hat{f} .

Let h represent the error induced by the projection,

$$f = \hat{f} + h, \quad (2.1.22)$$

and use $\|h\|^2$ as the error metric. Let the Gabor expansions under consideration have an $N \times M$ time \times frequency grid size; then there are MN points in any associated function (the data, window, Gabor coefficients and the Zak transforms). In all the following we assume the gaussian case with its zero at the center of the unit square. It is easy to see that

$$Z_{(h)}(p/N, q/M) = \begin{cases} 0 & ; (p, q) \neq (N/2, M/2) \\ Z_{(f)}(\frac{1}{2}, \frac{1}{2}) & ; (p, q) = (N/2, M/2) \end{cases} \quad (2.1.23)$$

and that

$$h(q/M + r) = \frac{1}{N} \sum_{p=0}^{N-1} Z_{(h)}(p/N, q/M) \exp(-j2\pi p / N) \quad (2.1.24)$$

$$= \begin{cases} 0; & q \neq M/2 \\ \frac{(-1)^r}{N} Z_{(f)}(\frac{1}{2}, \frac{1}{2}); & q = M/2 \end{cases} \quad (2.1.25)$$

One then readily finds that

$$\|h\|^2 = \frac{1}{N} |Z_{(f)}(\frac{1}{2}, \frac{1}{2})|^2. \quad (2.1.26)$$

We can use $\|h\|^2 / \|f\|^2$ as a measure of energy loss in the projection, and estimate its approximate value by noting that the loss is one out of MN coordinates; since the Zak transform is unitary, we can anticipate an energy loss on the order of $1/MN$.

Consider further that there are M unique alignments of f with the time grid in regard to Zak transformation, and there is no reason to imagine an a priori preference for any one of them. Define the set of time translates of f as $\{f^{(m)}\}$, $0 \leq m \leq M-1$, where we identify f with $f^{(0)}$, and $f^{(m)}$ is its m -th translate. Then in attempting to minimize the projection error, we are free to inspect all M translates of f and choose the one such that $|Z_{(f^{(m)})}(\frac{1}{2}, \frac{1}{2})|^2$ is least. The error associated with each translate is

$$\varepsilon_m^2 = \frac{|Z_{(f^{(m)})}(\frac{1}{2}, \frac{1}{2})|^2}{\sum_{p,q} |Z_{(f^{(m)})}(\frac{p}{N}, \frac{q}{M})|^2}, \quad (2.1.27)$$

and the minimum is

$$\varepsilon_m^2 = \min_{0 \leq m \leq M-1} \varepsilon_m^2. \quad (2.1.28)$$

We would like to estimate the savings in choosing to discard the smallest possible amount of energy. To do so requires some assumption about the signal being represented; let us take as an example a signal f consisting of statistically independent, real-valued, zero mean, unit variance gaussian noise variables. For any f it is true that

$$Z_f(\frac{1}{2}, \frac{1}{2}) = \sum_{k=0}^{N-1} (-1)^k f(k + \frac{1}{2}), \quad (2.1.29)$$

and therefore the new gaussian variable in (2.1.29) satisfies

$$\overline{Z_f(\frac{1}{2}, \frac{1}{2})} = 0 \quad (2.1.30)$$

and

$$\overline{|Z_f(\frac{1}{2}, \frac{1}{2})|^2} = \sum_{k=0}^{N-1} \overline{|f(k + \frac{1}{2})|^2} = N. \quad (2.1.31)$$

We would like to compare this average energy in the Zak transform of f to the expected energy in the Corresponding Zak of signal $f_{(\hat{m})}$, for which the energy loss is minimum. We can do this through the following theorem.

Theorem:

$$\overline{\left| Z_{f(a)}(\frac{1}{2}, \frac{1}{2}) \right|^2} \leq \frac{N}{\sqrt{M}}. \quad (2.1.32)$$

This theorem tells us that there is an expected savings proportional to the square root of the available number of shifts of the signal relative to the grid.

Proof: Let

$$S_m = Z_{f(m)}(\frac{1}{2}, \frac{1}{2}); \quad 0 \leq m \leq M-1 \quad (2.1.33)$$

and define

$$S = \min_m \{|S_m|\}. \quad (2.1.34)$$

We want to find the probability density of S and compute its second moment.

Clearly

$$\Pr\{S \geq X\} = \Pr\{(|S_0| > X) \cap (|S_1| > X) \cdots \cap (|S_{M-1}| > X)\} = \left[\Pr\{(|S_0| > X)\} \right]^M. \quad (2.1.35)$$

But

$$\Pr\{(|S_0| > X)\} = \Pr\{(S_0 > X) \cup (S_0 < -X)\} = 2Q\left(\frac{X}{\sqrt{N}}\right), \quad (2.1.36)$$

where $Q(\cdot)$ is the normal probability tail integral

$$Q(x) = \int_x^\infty du p(u) = \frac{1}{\sqrt{2\pi}} \int_x^\infty du \exp(-u^2/2), \quad (2.1.37)$$

and $p(\cdot)$ is the gaussian density. We observe that

$$\frac{d}{dx} Q(x) = -p(x). \quad (2.1.38)$$

The probability distribution (integral of the density) of s is then given by

$$\Pr\{(S < X)\} = 1 - \left[2Q\left(\frac{X}{\sqrt{N}}\right) \right]^M, \quad (2.1.39)$$

and its density is found by differentiation:

$$q(x) = \frac{d}{dx} \Pr\{(S < X)\} = \frac{2M}{\sqrt{N}} p\left(\frac{x}{\sqrt{N}}\right) \left[2Q\left(\frac{X}{\sqrt{N}}\right) \right]^{M-1}; x \geq 0. \quad (2.1.40)$$

Then

$$\begin{aligned} \overline{S^2} &= \int_{-\infty}^{\infty} du x^2 q(x) = \frac{2M}{\sqrt{N}} \int_0^{\infty} dx x^2 p\left(\frac{x}{\sqrt{N}}\right) \left[2Q\left(\frac{X}{\sqrt{N}}\right) \right]^{M-1} \\ &= 2MN \int_0^{\infty} dv v^2 p(v) [2Q(v)]^{M-1}. \end{aligned} \quad (2.1.41)$$

We see that (2.1.41) yields a value of N for the case $M = 1$, which checks with (2.1.31). Although we cannot analytically carry out the integration in (2.1.41) to get a closed form exact value, we can get a tight upper bound by using the familiar inequality

$$Q(v) \leq \frac{1}{2} \exp(-v^2/2); v \geq 0 \quad (2.1.42)$$

as follows:

$$\begin{aligned} \overline{S^2} &= 2MN \int_0^{\infty} dv v^2 p(v) [2Q(v)]^{M-1} \leq 2MN \int_0^{\infty} \frac{dv}{\sqrt{2\pi}} v^2 \exp(-Mv^2/2) \\ &= \frac{2N}{\sqrt{M}} \int_0^{\infty} \frac{dy}{\sqrt{2\pi}} y^2 \exp(-y^2/2) = \frac{N}{\sqrt{M}}. \quad \text{Q.E.D.} \end{aligned} \quad (2.1.43)$$

Our interpretation of this result is as follows. If we arbitrarily position our signal with respect to the grid, we can expect to lose the energy of one of our MN coordinates. If we replace the numerator and denominator in (2.1.27) with their expected values, we find

$$\epsilon_m^2 = \frac{|Z_{(f^{(m)})}(\frac{1}{2}, \frac{1}{2})|^2}{\sum_{p,q} |Z_{(f^{(m)})}(p/N, q/M)|^2} = \frac{1}{MN}. \quad (2.1.44)$$


Choosing the signal placement wisely lets us achieve

$$\epsilon_m^2 = \frac{1}{\sqrt{M}} \left(\frac{1}{MN} \right), \quad (2.1.45)$$

i.e., the expected fraction of energy loss is smaller than the fraction of coordinates discarded. Thus from an energy viewpoint, representation by a discrete Gabor transform using a window having a single zero may be quite tolerable!

There is another perspective to put on this sequence of observations. We have selected for our signal a stationary process, and have applied to it a time-frequency transform. Conventional wisdom tells us the utility of time-frequency techniques is that they capture nonstationary behavior. In one sense then, we might not expect any benefit from the analysis, and one could argue such a case, saying "our window selection was so inept that the basis functions were incomplete and we lost some signal; only by a clever trick did we minimize that loss."

In support of this argument, notice that the expected energy loss increases as N increases, i.e., as the time resolution of the transform increases. This may be considered to be the penalty associated with going from Fourier-type representations, which have no time resolution, to a mixed domain picture when the phenomenon under study is stationary. Now suppose we give our signal a second-order statistical cyclostationarity by assuming the variance to be a periodic function with a period such that several cycles are captured within the MN data points. In the period is in fact \tilde{M} , we would anticipate that a Gabor representation using \tilde{M} frequency points and $\tilde{N} = MN / \tilde{M}$ time points to be well matched to the signal. If the variance had a pattern such as that shown in Fig. 2.1-1, one can see that by picking the grid such that the time coordinate of the zero of the window matches the minimum variance point, one would in fact expect to lose little and do better, on the average, than the $1/\sqrt{M}$ improvement found for the stationary case. Here we are exploiting the time-varying behavior of the Gabor series to match the problem, and are finding some good rewards.

Fig. 2.1-1. Graph of the periodic variance of σ .

Observing all that has been done above, one can conclude that it is not actually necessary to sacrifice *any* of the signal energy, if we will liberalize our representation slightly. The space into which we have projected the signal has an orthogonal complement spanned by, in this case, a single basis function. If we add this function to the $(MN - 1)$ others generated by the time/frequency translates of the Gabor window, we recover a complete basis. It differs from a strict Gabor basis in that one member fails to exhibit the translation property. This gives us another way to see the penalty for using the window with a Zak zero; it causes us to abandon part of the nice structure we valued in going to the Gabor series in the first place. However, the deviation can be considered minimal. In fact, the added basis function depends only on one of each M points in the data; this function is just $g = h / Z_f(\frac{1}{2}, \frac{1}{2})$, where h is given by (2.1.22);

$$g = \begin{cases} 0; & q \neq M/2 \\ \frac{(-1)^q}{N}; & q = M/2 \end{cases} \quad (2.1.46)$$

Observe g is zero except for one point in M .

The above viewpoint says that we need not abandon a window function that resembles our data quite well simply because of its Zak zeros. The options are to remain within the strict Gabor structure and sacrifice a small amount of signal energy, or to enlarge the representation and capture everything. One of these approaches may well suit a wide variety of problems.

2.1.3 Relations between Matrix Stability and Signal Duration

The Matrix Gabor representation as described in [BALA2] has some stability problems associated with it that are mentioned in the original paper, such as noninvertibility of the window if it possesses an analytical (or computational) zero in the first Gabor time slice, as well as other problems that are mentioned in the previous section and were investigated in [BALA1], a copy of which is included in Appendix A. We will now describe some preliminary results which lead us to believe that one may have some control over the stability of the method by controlling the signal length or the number of Gabor time points, depending on the nature of the signal that is to be analyzed. For windows whose values in the first Gabor time slice are not zero, but have a big dynamic range and with the smallest value being much less than one, the method is also unstable as can be seen from the following. Let the linear system to be solved be defined as

$$(WE)x = b \quad (2.1.47)$$

where W is the $MN \times MN$ (block) matrix containing the window values, E is the $MN \times MN$ (diagonal) Fourier rotation matrix, x is the $MN \times 1$ vector of (ordered) unknown coefficients, and b is the $MN \times 1$ vector of signal data points, M is the number of Gabor frequency points and N is the number of Gabor time points. We can now write the vector x of coefficients as

$$x = E^{-1}Bb, \quad B = W^{-1} \quad (2.1.48)$$

Let us now write W and B in their block form in order to see their representation more clearly and to see how B will depend on the structure of W . Note that because of the ordering each of the blocks of B and W will be $M \times M$, and there will be $N \times N$ of them in each matrix (for a full derivation see [BALA2]).

$$W = \begin{bmatrix} W_0 & & & & & \\ W_1 & W_0 & & & & \\ W_2 & W_1 & W_0 & & & \\ \vdots & \ddots & \ddots & \ddots & & \\ W_{N-2} & & W_2 & W_1 & W_0 & \\ W_{N-1} & W_{N-2} & \cdots & W_2 & W_1 & W_0 \end{bmatrix} \quad (2.1.49)$$

$$B = \begin{bmatrix} B_0 & & & & & \\ B_1 & B_0 & & & & \\ B_2 & B_1 & B_0 & & & \\ \vdots & \ddots & \ddots & \ddots & & \\ B_{N-2} & & B_2 & B_1 & B_0 & \\ B_{N-1} & B_{N-2} & \cdots & B_2 & B_1 & B_0 \end{bmatrix} \quad (2.1.50)$$

and

$$\begin{aligned} B_0 &= -W_0^{-1} \\ B_1 &= -W_0^{-1}(W_1 W_0^{-1}) \\ B_2 &= -W_0^{-1}(W_2 W_0^{-1} + W_1 B_1) \\ B_3 &= -W_0^{-1}(W_3 W_0^{-1} + W_2 B_1 + W_1 B_2) \\ &\vdots \\ B_l &= -W_0^{-1}(W_l W_0^{-1} + W_{l-1} B_1 + \cdots + W_1 B_{l-1}); \quad \forall l \in [2, N-1] \end{aligned} \quad (2.1.51)$$

The extent of our analysis to date has been to look at the simplest nontrivial case which occurs when the window lasts only over the first two Gabor time slices. This being the case, the zeroth and the first W are nonzero, and the above formula for the biorthogonal matrix becomes

$$\begin{aligned} B_0 &= -W_0^{-1} \\ B_1 &= -W_0^{-1}(W_1 W_0^{-1}) \\ B_2 &= -W_0^{-1}(W_1 B_1) = W_0^{-1}(W_1 W_0^{-1})^2 \\ B_3 &= -W_0^{-1}(W_1 B_2) = -W_0^{-1}(W_1 W_0^{-1})^3 \\ &\vdots \\ B_l &= -W_0^{-1}(W_1 B_{l-1}) = (-1)^l W_0^{-1}(W_1 W_0^{-1})^l; \quad \forall l \in [2, N-1] \end{aligned} \quad (2.1.52)$$

from which we notice two important things. The first is that even though the window only has support over two time slices, the biorthogonal has support over all the time slices, therefore, the biorthogonal matrix will be full. More apropos to this discussion, we notice that each successive block of the biorthogonal grows by a power of the inverse of the window entries in the zeroth time slice multiplied by the corresponding window entries in the first time slice. This is due to the diagonal nature of each of the blocks of W and B . This implies that the stability of the method depends on containing that growth rate to reasonable bounds.

The above discussion implies that one now has two possible ways of controlling the stability of the method depending on the nature of the signal. If the signal of interest is such that it can tolerate being split up into pieces, then one can choose the desired number of frequencies M and then make the length of the signal be such that the number of time points N keeps the power

$$B_{N-1} = (-1)^{N-1} W_0^{-1} (W_1 W_0^{-1})^{N-1}. \quad (2.1.53)$$

well behaved and bounded. One can then apply the transform to each of the pieces separately. If on the other hand, the signal in question does not tolerate being split up, but we have some freedom in choosing the number of frequency points, ie, the window is nonzero over the increased length of the Gabor time slice, then one can increase the number of frequency points thereby decreasing the number of time points (for a fixed signal length) and keeping the above expression well behaved.

The computation of the biorthogonal matrix entries gets more involved as the number of Gabor time points that support the window increases, and this analysis has not yet been performed for the general case. When the analysis is complete, we will submit a paper for publication on the obtained results. As a closing note we remark that in GSPS, all the aforementioned analysis can be done a-priori for windows of all lengths of interest, and the information will be readily available when the signal is chosen, since the analysis is completely independent of the signal.

2.2 ALGORITHMS

2.2.1 Oversampling with the Gabor Transform

In order to increase the capability of the GSPS workstation to allow us to analyze different types of signals, we decided to include the capability to perform oversampling. A brief description of the algorithm is given here, and for a full detailed explanation the reader is referred to [WEXL]. The formulation of the oversampling algorithm is as follows: Let P be the total number of points in the signal data set, M and N be the number of desired frequency and time points respectively, and M' and N' be two auxiliary positive integers satisfying the constraint that

$$P = MN' + M'N \quad (2.2.1)$$

Under these conditions, the Gabor transform can be written as

$$s(k) = \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} a_{m,n} \tilde{h}(k - mN') \exp(2i\pi nk / N) \quad (2.2.2)$$

$$a_{m,n} = \sum_{k=0}^{P-1} s(k) \tilde{b}^*(k - mN') \exp(2i\pi nk / N) \quad (2.2.3)$$

together with the condition that

$$\sum_{k=0}^{P-1} h^*(k + mN) \tilde{b}(k) \exp(2i\pi nk / N') = (P / MN) \delta_m \delta_n. \quad (2.2.4)$$

At this stage, the last condition can be rewritten in matrix form. in order to solve for the biorthogonal function \tilde{b} in a similar way as the matrix formulation of the Gabor transform. The $M' \times M$ block matrix of window values, each square block being $N' \times N'$, can be written as

$$(WE)b = r \quad (2.2.5)$$

with

$$W = \begin{bmatrix} W_0 & W_N & W_{2N} & \cdots & W_{(M-2)N} & W_{(M-1)N} \\ W_M & W_{M+N} & W_{M+2N} & \cdots & W_{M+(M-2)N} & W_{M+(M-1)N} \\ W_{2M} & W_{2M+N} & W_{2M+2N} & \cdots & W_{2M+(M-2)N} & W_{2M+(M-1)N} \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ W_{(N-1)M} & W_{(N-1)M+N} & W_{(N-1)M+2N} & \cdots & W_{(N-1)M+(M-2)N} & W_{(N-1)M+(M-1)N} \end{bmatrix}, \quad (2.2.6)$$

with each W_l given by

$$W_l = \begin{bmatrix} h_l & & & & \\ & h_{l+1} & & & \\ & & h_{l+2} & & \\ & & & \ddots & \\ & & & & h_{l+(N-2)} \\ & & & & & h_{l+(N-1)} \end{bmatrix}, \quad (2.2.7)$$

for a grand total of $M' N' \times P$ entries, premultiplied by the $M' \times M'$ block matrix E_0 of Fourier rotations of size N' given by

$$E = \begin{bmatrix} E_0 & & & & \\ & E_0 & & & \\ & & E_0 & & \\ & & & \ddots & \\ & & & & E_0 \\ & & & & & E_0 \end{bmatrix} \quad (2.2.8)$$

and the $P \times 1$ unknown biorthogonal and the $M' N' \times 1$ right hand side vectors b and r are respectively,

$$\begin{aligned} b &= [b(0) \ b(1) \ b(2) \ \dots \ b(P-2) \ b(P-1)]^T \\ r &= [P/MN \ 0 \ 0 \ \dots \ 0 \ 0]^T \end{aligned} \quad (2.2.9)$$

The biorthogonal vector can be now calculated after choosing a way to 'invert' the non-square matrix associated with the window function. We chose to implement a generalized inverse or energy minimization method in the GSPS workstation with some slight computational modifications from the referenced paper. We noticed that since the inverse is given by

$$\begin{aligned} b &= (EW)^T ((EW)(EW)^T)^{-1} r \\ &= W^T (WW^T)^{-1} \hat{r} \\ \hat{r} &= E^{-1} r \end{aligned} \quad (2.2.10)$$

the last computation can be performed by hand and it results in a reasonable computational savings. The new right hand side vector becomes

$$\begin{aligned} r &= [r_0 \ 0 \ 0 \ \dots \ 0 \ 0]^T \\ r_0 &= [P/MN \ P/MN \ P/MN \ \dots \ P/MN]^T \end{aligned} \quad (2.2.11)$$

with each one of the blocks being of size $N' \times 1$. The biorthogonal is, at this stage, fully available and the Gabor coefficient map can be computed with the aid of (2.2.3). Reconstruction of the signal (with or without postprocessing) can now be also calculated with the aid of (2.2.2). A better description on how to perform these operations within GSPS and a description of which parameters can be chosen by the user will be described in Section 3.3.

CHAPTER 3

SOFTWARE DEVELOPMENT

3.1 SUMMARY OF SECOND YEAR EFFORT

Before discussing the software effort expanded this third and last year of the contract, we will briefly summarize the work that was performed during the second year under contract.

3.1.1 Phase II Code

During the second year of effort, numerous additions and enhancements were incorporated into the Gabor Signal Processing System software, which from here on we will call GSPS. Before moving on to summarize the Phase II code, we will begin by describing the philosophical change that took place towards the end of the first year of the contract. As described in the First Annual Technical Report, the initial Gabor transform software prototype was implemented on an IBM compatible PC and it was command line driven, ie, the user was prompted to enter a number corresponding to a signal, window, etc., the Gabor processing was performed, and the results were then written to files. Unfortunately, in order to change parameters, one had to either restart the code or hardwire different values into the program (such as total number of points) and recompile. The cumbersomeness of doing this, together with memory limitations and speed considerations, forced us to migrate to a different (better) platform and to create a user friendly graphical user interface that could accomodate all the power and flexibility that we wanted to incorporate into the GSPS. Bearing that in mind, at the beginning of year two we selected a SUN Microsystems SPARCStation1+[®] as the computational platform, C as the scientific programming language, and X11 together with the XView toolkit as the graphical user interface programming language to be used to develop the GSPS software that is currently available, and which we have been using 'in house' for a period of well over one year.

3.1.1.1 Functionality

The organization of the program has become more clear to the user by the development of a graphical user interface that is layed out in such a way that the user can follow a sequence of major categories, represented on the screen as buttons, and progresively move down the logical sequence to the desired function to be applied. This is accomplished in the following manner: the major categories represented by buttons are:

Signals
Windows
Noise
Method
Clipping
Reconstruction
Iterate
Utility
Coefficients
Options
Output
Exit GSPS ?

Each of the buttons is the owner of a menu which contains menu items, which in turn may contain (nested) submenus. The last entry of a path through a menu is either a command which is executed immediately or, when appropriate, a command which contains a dialog box in which the user can either enter, change, or choose selected default parameters. These parameters are then updated in the main program and the command is executed. If the action that is selected by the user results in a graphical representation of data, a window is automatically opened and the data is displayed. Since the real estate on a computer screen is limited, all of the windows that are opened by GSPS are multifunctional, ie, the same graphical window that is used to display the analysis window is used to display the function that is biorthogonal to said analysis window.

Keeping that in mind, we can now illustrate the high level execution of the program by means of an example, while a more detailed example including the features incorporated in the third year will be included in the User Manual (Section 3.3). The user starts execution of GSPS by opening their favorite X11-based window manager on the screen, and typing *gps* from one of the text I/O windows, or alternatively, opening a file manager and double clicking on the *gps* application. This results in the display of a Control panel containing the (deactivated except for Exit GSPS) buttons, and a Disclaimer panel containing the proprietary information, disclaimer, software version, and two buttons marked "Continue" and "Quit". Selecting the Quit button exits GSPS with no action being taken, while selecting the Continue button erases the Disclaimer panel and activates the buttons in the Control panel. We are now ready to start processing signals. Selecting the "Signals" button with the Right Mouse Button (RMB)displays a menu containing different signal choices right under the Signal button. Selection of a signal choice with the Left Mouse Button (LMB) opens up a dialog box which allows the user to either choose

default values for the signal by selecting the "Defaults" button in the dialog box, or enter the user defineable parameters manually. The action can now be either aborted by selecting the "Cancel" button, or accepted by selecting the "Ok" button. Accepting the choice results in the appearance of a display window which contains the selected signal. One can now proceed to the selection of the analysis window by selecting the "Window" button with the RMB. This results in the display of a menu that is functionally identical to the "Signals" menu and, in fact, contains the same entries. Selecting the "Ok" button from the Dialog box results in the appearance of a second display window that contains the graphical representation of the analysis window. Once the signal and window have been chosen, the user can proceed to the choice of the method that will be used to do the processing by choosing a method from the menu associated with the "Method" button. This can be accomplished by selecting the "Method" button with the RMB, and choosing a method with the LMB results in a third display window being open, and it contains the appropriate coefficients corresponding to the selected method. It is noteworthy to mention that the Coefficients window automatically displays both 2-D and 3-D coefficient sets depending on the representation that is warranted by the given method. At this stage the user can choose to repeat the portion of the experiment described so far by choosing different signals, windows, methods, or any combination of the above, or proceed to the reconstruction of the signal. If one opts for the reconstruction route, there is an option to first perform simple thresholding on the coefficient set by selecting the "Clipping" button with the RMB, entering the desired thresholding levels, and removing the coefficients below, above, or between user defined upper and lower bounds. If it turns out that it is not desirable to perform any thresholding, one can go directly to one of the reconstruction routines which behave analogously to the method options in the "Method" button. The reconstructed signal is now displayed in the Signal window and, if the reconstruction method is the same as the transform method, the coefficient set has not been clipped, and the method was stable, it will be the same as the original signal except for roundoff error. The L_2 distance between the original signal and the reconstructed signal is reported next to the origin in the Signal display window

3.1.2 Capabilities

During the second year of effort, numerous additions and enhancements were made to the tool that was available at the end of the first year, the most noteworthy of which being that the user could change initialization parameters like signal duration and time-frequency splitting of the grid, as well as other parameters, without having to recompile the program. In fact, the only time that recompilation is needed is when a new feature is added to one of the menus as a *menu*

item. After the addition is debugged, all the parameters that can be controlled by the user (if any) will appear in a dialog box when the selection is made.

Another major improvement was the feel of the tool. The user could now input parameters of interest in the dialog boxes, and choosing the button labeled 'Ok' made the program accept all the parameters at once and continue with the execution. If the wrong parameter was input, the user could move the pointer over to the appropriate field and change it before selection of the 'Ok' button if the mistake was noticed in time or, in the case that it was not, the same (or a different) menu item could be chosen, the correct parameter input and the 'Ok' button selected. In the previous version of the software, clerical errors would result in the user having to restart the program from step one. The menu approach also allowed the user to pick up at any point on the list of buttons which logically preceded the last step chosen, and after an initial run through the four logically sequential steps, ie, after going through the process of choosing the signal, window, method and reconstruction, one could reprocess from any of the intermediary steps. We will illustrate by saying that after the user had chosen the signal, window, and method, the signal, window or method could be chosen again, as well as the associated utilities, without having to either complete the sequence by choosing the reconstruction option or starting out with a new copy of the same (or different) signal. The logical flow of the program is illustrated in Fig. 3.1-1 below.

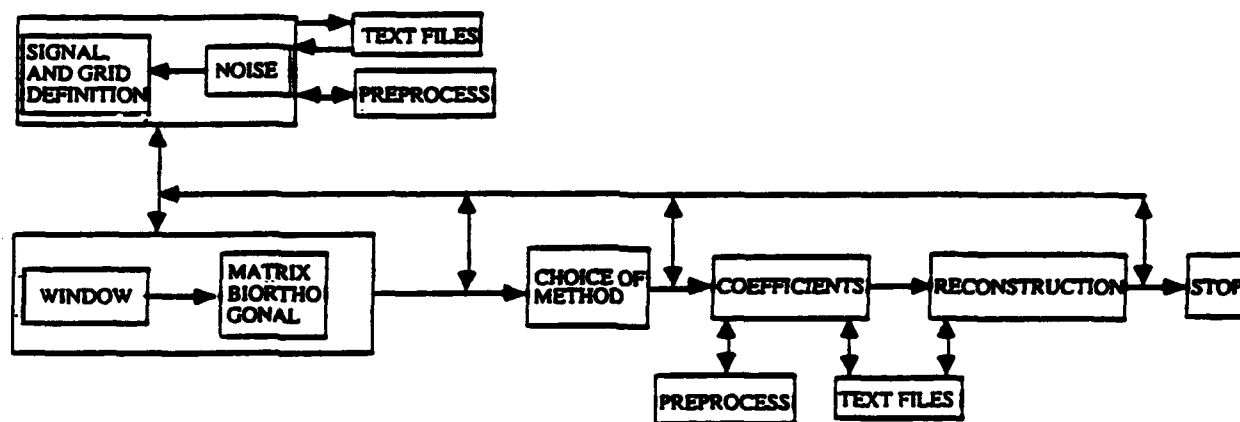


Fig. 3.1-1. Logical flow of the GPS package at the end of the second year

Note that even though the user could choose a different window or method after having gone through the sequence of events once, the grid definition was at the signal level, therefore, if one wanted to redefine the grid shape and size, one would have to go back to the signal choice and redefine the signal in such a way that it reflected the desired configuration. We will see in the next section that this limitation no longer exists due to the ability of the window to redefine

grids, and the added capability to redefine grid configuration from a menu item's dialog box under the Utility button.

During the second year of effort we also included the method resulting from the matrix reformulation of the transform [BALA] and numerous utilities to preprocess the signal, such as zooming, shifting, saving to memory, etc.. We also reported various values of interest to the appropriate window after certain actions were taken, such as the length of the maximum coefficient and the dimension of the signal set [ORR?]. For a full description of the enhancements introduced during the second year the reader is referred to [AAE1].

3.2 THIRD YEAR EFFORT

3.2.1 New Facilities

As reported in the previous technical report, at the end of year one it was decided to improve our computational facilities in order to be able to process larger data sets, and to speed up the processing time. As it turns out, the upgrade that was made to the system at that time has resulted in more than adequate computational facilities to perform the tasks associated with this project and, therefore, no new facilities were needed during the third year of effort. For a full list of facilities, the reader is referred to [AAEC] and [AAE1].

3.2.2 Phase III Code

Phase III software development has centered around broadening the applicability of GSPS to a wider range of problems by improving its ability to interoperate with other analysis packages, and through the addition of new diagnostic displays. In addition, we have implemented two new versions of the Gabor transform with slightly different capabilities than those present in the Phase II code. To facilitate the addition of these features to GSPS, several key portions of the Phase II code have been restructured for added flexibility, notably in the file I/O facilities and in the handling of signal buffers.

3.2.2.1 New Display Components

map view

A "map" mode has been added to the coefficient display to supplement the spike plot display from Phase II. In this mode, the coefficient matrix is displayed as a grid of rectangular color patches whose colors correspond to the magnitude of the coefficient at that coordinate, relative to the maximum coordinate magnitude. Coefficient values below a compiled-in epsilon

value are suppressed, and are displayed in the background color. A color key is provided to the lower left of the grid. Selection of the spike plot and map modes is accomplished using the "map" button located in the control area above the color key.

grid button

A button labeled "grid" now appears in the control areas of the signal, window, and coefficient displays. In the signal and window displays, this option controls the display of the hash marks which indicate the endpoints of the Gabor timeslices. In the coefficient display, this option works in conjunction with the map button to control the display of reference lines in the spike plot mode, or to mark the boundaries of each color patch in the map mode.

The RMS duration and RMS bandwidth are now computed for both the signal and window, and are displayed in the control areas of their respective windows.

A number of new signal options are available:

(interaction changes: new signal and window types)

Get Fourier

null signal

(interaction changes: reading data files)

getfileinfo, matlab

(interaction changes: new transforms)

matrix2; oversampling

(interaction changes: piecewise forward and reverse transforms)

the "compute" routines

biorthogonals and windows: matrix and zak

forward and reverse zak

(interaction changes:)

(changes to utilities)

One significant departure from the Phase II code is that the buffers containing the intermediate components of the Gabor transforms (i.e. the biorthogonal and Zak buffers) may now be loaded from a file and explicitly manipulated. This introduces a synchronization problem as these buffers were implicitly linked to other buffers when their contents were automatically generated. A set of global flags have been added to indicate the validity of these

buffers, and should be checked before the corresponding buffers are to be used. Likewise, these flags should be set appropriately whenever any of the buffers is modified.

3.2.3 Warnings and Suggestions

The following is a list of known caveats for potential users.

After the tool has been running for a long time, it may crash.

Since you can't 'edit data' in the 'window' window, any Gabor window preprocessing can be done in the signal window, copied to memory, and then read in the window window.

The number of points in a signal or window should be an integer power of 2.

Although the coefficient display is cleared whenever it becomes marked as invalid, a refresh call to that window may cause the previous contents to be redisplayed.

A number of conventions must be observed when Matlab data files are to be imported into GSPS. Several variable names are reserved for use by GSPS [tmpM tmpN tmpRate tmpType tmpMethod Comment]. Any stored data elements which do not correspond to one of these names will be interpreted as the data with which the buffer will be filled. As a result, you should not save any other elements in this file except for the signal itself, which may be saved with an arbitrary name. If the file importing routine encounters an unknown element following the signal, it will attempt to load that element's contents instead.

GSPS was developed in the OpenWindows™ 2 environment, and is designed primarily for use in a color environment. Although the application will run successfully with a monochrome display, various text fields will appear crowded as the monochrome XView™ 2.x library uses a different font entirely. GSPS will compile and run with XView 3.x libraries as well, but will exhibit this behavior on both color and monochrome displays.

3.2.4 Scripts

Although the Phase II version of GSPS included a simple command line interface, it has since proven inadequate to the expanded capabilities of the Phase III code. As GSPS has become more general in its purpose, so has the need for a more generic command line interface grown. The new interface syntax is designed to accommodate additional transforms and also makes provisions for transform-specific parameters.

gsps -s <fname> -w <fname> -x{2lmlolz} -o <fname>

gsps -s <fname> -w <fname> -m <M> -n <N> -x{2lmlolz} -o <fname>

where:

-s <fname> is a path to the file containing the signal

-w <fname> is a path to the file containing the window

-x specifies the transform to be used:

-x2 indicates the Matrix2 method

-xm indicates the Matrix method

-xo indicates the oversampling method

-xz indicates the Zak method

-o <fname> contains the name of the file to which the coefficients should be written

3.3 USERS GUIDE

3.3.1 Procedure Descriptions

The following new files and routines are present:

new files: get_four_coef.c getfileinfo.c loadmat.c null_sig.c over_gabor.c savemat.c tbprod.c
utilmat.c

new routines: graphics.c init.c loadfile.c main.c savefile.c xforms.c

In the event that a signal, window, or coefficient matrix of exceptionally small magnitude must be displayed, the plotting routines now impose a non-zero scale value...

Example: how to add a new signal type

Although new signals are typically generated externally and imported through the file I/O interface, it is frequently desirable to hardcode new signal types which are to be used frequently with minor variations, even though these signals will only be accessible from the graphical user interface. This procedure typically requires four steps: coding the signal generator, coding the

user interface components, adding calls to the main initialization and menu generation routines, and adding build dependencies to the compilation control configuration file. For clarity, it is generally desirable to name the source code files after the signal being generated; for instance, a triangular pulse generator might have the files `tri_pulse.c` and optionally, `tri_pulse.h` associated with it. Because GSPS lacks a generic mechanism for manipulating signal generation parameters outside of the GUI, these variables should remain opaque to the remainder of the program and may be defined as static types local to the signal's source module.

The signal generation routine is often adapted from another piece of standalone software, so typically it is passed a desired number of points, a sample rate, and whether a real or complex-valued signal is desired, and typically it returns a buffer of the appropriate size and type. This may then be called from a callback routine which is shared by the signal and window menu items, which is then responsible for copying the generated signal into the correct buffer.

The GUI supporting code consists of an initialization routine which creates the GUI objects, and a set of callback routines which those objects will dispatch as various buttons and widgets are manipulated. The construction of these callbacks is beyond the scope of this document, but the file `tri_pulse.c` is recommended as a template.

Modifications should be made to the file `main.c` in four places. First, external references to the routines should be added near the top of the file. Next, menu items should be created under both the Signal and Window menus. Finally, a call to the initialization routine should be added near the end of the file, but before the main loop is invoked.

3.3.2 Program Flow

During the third year effort, much functionality was added to the GSPS software, but the logical flow of the program has changed very little since the end of the previous year. The main logical change has been due to the inclusion of a menu item called *grid, info* to the Utility button. This menu item allows the user to redistribute the time and frequency points on the grid at any time after the signal buffer has been written to at least once has been added to the utility button. therefore, the user no longer needs to choose a new (or the same) signal to run different grid configurations. The user still needs to generate a new signal buffer if the signal duration is to be changed. One other change worth mentioning is that under the Utility button, a Compute menu item has been added which allows the user to compute intermediate steps in the Gabor transform computation, for instance, the Zak transforms and the window biorthogonal. This means that the user no longer has to perform the full transform computation if the only thing of

interest at any particular time is, say, the biorthogonals to the window that result from a set of grid spacings. For big signals, this can prove to be a considerable time savings

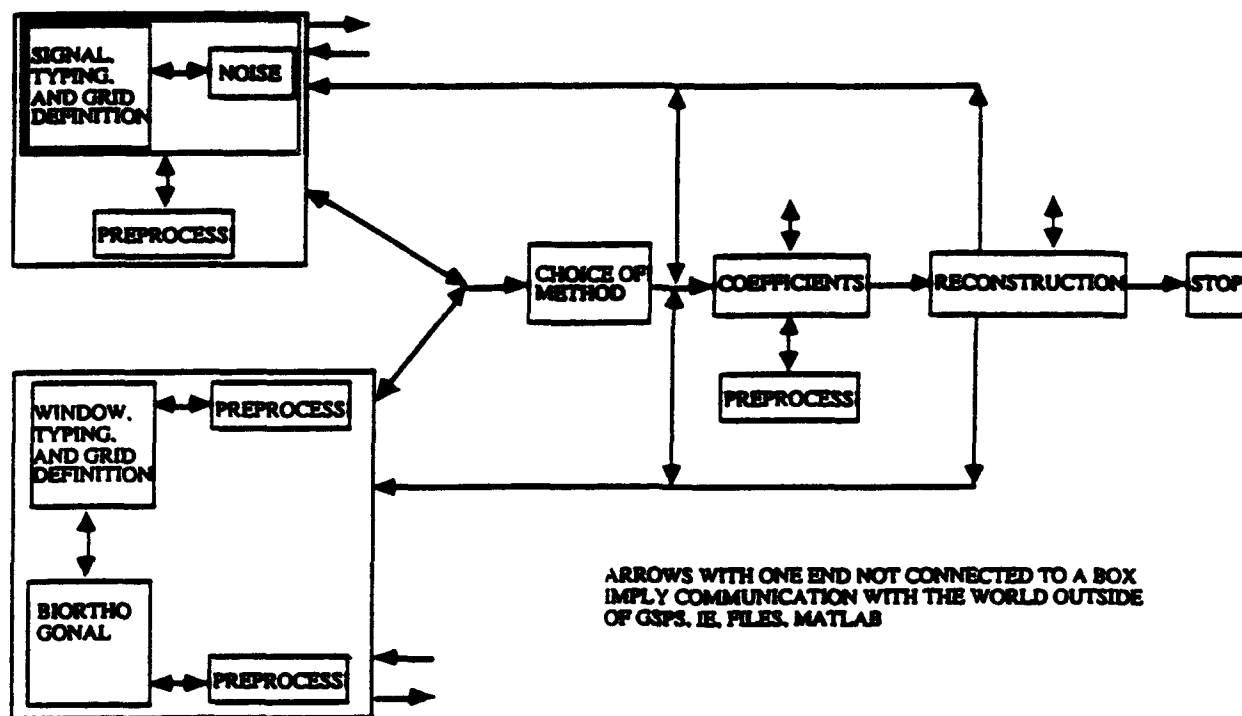


Fig. 3.3-1. Illustration of program flow.

CHAPTER 4

APPLICATIONS RESEARCH

4.1 FAULT IDENTIFICATION IN FEEDBACK CONTROL CIRCUITRY

This section describes the experiments performed to investigate the application of the Gabor representation for solving problems of emitter identification in communications. In these experiments, test data was used to investigate the characterization of signal features, specifically transient signals, with the Gabor transform. These test data features were used to characterize signals as belonging to specific classes of emitters as an exercise in determining utility of Gabor-based methods for this problem

The problem investigated in these experiments was the mapping of signals of interest to transmitting and receiving equipment identity. The utility of using Gabor-based techniques was studied and tested with a set of signal data from several transmitter-receiver pairs. Each pair was considered to generate a class of signals. The Gabor techniques were applied to the data to characterize each signal as a *member of one of these classes of signals*. A dimension measure of the Gabor coefficient set, developed by AAEC under a separate DARPA contract [AAEC2], was the main discriminant tested to separate the signals into classes.

4.1.1 Summary of Previous Work (Greenbelt Facility)

4.1.1.1 Data Sets

Two sets of signal data were used during this effort. The first database consisted of 186 files containing approximately 65 thousand time samples each, sampled from analog signals at a 20 Khz rate. Of the 186 files, only 47 files contained signals of interest. Two types of transient features were found in these 47 files, and 56 occurrences of these features were chosen for the experiments. The two types of transient features were designated the small and normal size signals. The length of these features of interest ranged from 256 to 2048 time samples.

A second database consisted of 1944 files, containing signal data from many different transmitter-receiver pairs. The length of these files ranged from approximately 1000 to 16500 time samples. In this data set, a signal of interest contained about 200 time samples.

4.1.1.2 Experiments

The technical approach used in this investigation was to visually examine the original signal data and look for features in the signals which appeared to be common within each class of signals. Once these types of features were identified, a portion of the data containing the feature was extracted to be used as both a signal and as a Gabor window function. Data-derived window functions were then applied to the other signals using the Gabor transform. The resultant dimension of the set of Gabor coefficients was minimized through a time alignment procedure, and this minimum dimension was used to discriminate between classes. Additional discriminants were used as appropriate. These included the maximum amplitude of the Gabor coefficients, and the reconstruction error due to "clipping" of the Gabor coefficients, which is defined as deletion of all coefficients of magnitude smaller than a percent of the largest magnitude coefficient. These additional measures improved the classification process.

Three types of experiments were performed: 1) Normal vs small signal characterization, 2) Small signal discrimination, and 3) Signal classification with averaged windows. Since signals of interest in the data come in two classes, the first task was the discrimination between these, addressed by the first experiment. Given success at this first stage, discrimination within a class was the remaining key factor. Experiments 2 and 3 were devoted to discrimination among small signals. Discrimination among the normal size signals was not emphasized.

4.1.1.3 Results and Observations

Even though only a finite number of test signals and six window functions were used to analyze and characterize the signals of interest, use of the minimum Gabor dimension value and the maximum Gabor amplitude showed promising results. The use of averaged windows was found to be a useful additional method of discrimination. For a complete description of these experiments, refer to the Second Annual Technical Report (AAEC2).

4.1.2 Summary of Previous Work (Waltham Facility)

During the previous year, a MIPR subprogram was conducted under the Gabor program to address the problem of identifying specific emitters from a particular class of interest to the MIPR sponsor. A number of very large (megabyte-size) signals were provided by the MIPR sponsor. Each typically contained several transient signals of roughly a thousand samples each. The sponsor provided the locations of a few of these transients. The remainder (which comprised the majority of the available transients) were extracted by hand by plotting the million-sample

signals as 1000 x 1000 "images" and manually marking interesting segments with an interactive, cursor-based utility which we developed for this purpose. Later in the program, a morphological algorithm for automatically detecting and locating transients was developed and partially tested. On the MIPR sponsor's request, work on automatic detection and extraction was terminated so that the efforts could be concentrated on emitter classification.

Morphological filtering comprised the core of the classification techniques developed and tested, particularly for signal conditioning and feature extraction. The development of reliable, quickly computable, morphology-based discriminants and means for utilizing them was continued into the third year of the Gabor program, and is discussed in the next section.

4.1.3 Year 3 Effort (Waltham Facility)

A second MIPR was added to the Gabor program to continue work on the specific emitter identification problem. The primary technical objective was to demonstrate the effectiveness and efficiency of morphological and other modern signal processing techniques for specific emitter classification using transient signals provided by sponsor. The four key steps performed towards this end were the preprocessing of selected transients, feature extraction, classifier design, and performance evaluation. Morphological techniques were applied during the first two steps, and were found to have the most benefit for feature extraction, providing a compact, easily computed representation of the transients.

Two data sets were provided by the MIPR sponsor during Year 3, hereafter referenced as the old and new data sets, respectively. The old data set comprised 23 moderately long signals, typically containing 1000 samples each. Since our main thrust was classification, *i.e.*, not transient detection and location, we manually extracted a single segment, typically 100 samples long, from each signal. Although the Year 3 "old data set" signals were much smaller than those from the previous year, they still required manual transient extraction since they also contained other signal components which the sponsor specifically instructed us to ignore. The signatures taken from the old data set were fairly clean. The ground truth provided with them divided them into six classes, and appeared reasonable based on both visual and computer-aided analysis.

The new data set, also provided by the (MIPR) sponsor, was much larger, consisting of 1944 moderately long signals of roughly 1000 samples each. 180 segments, typically 100 samples long, were manually extracted from the collection. In contrast with the old signatures, many of the new ones had a very noise-like appearance, probably due to distortion, dispersion, widely varying band-limiting, and other effects induced by the transmission channel through which the

signals were received. Furthermore, it was difficult to relate the ground truth, which segregated the signatures into 11 classes, to the observed waveform characteristics.

Similar experiments were conducted on both of the aforementioned data sets. A high degree of success was achieved with the old data: all but one of the 23 signals were classified correctly. The lessons learned in both defining the architecture of the various classifier components and fine-tuning the associated algorithmic parameters were then applied to the new data. Although a number of different feature-space representations were tested and a moderate degree of success was obtained, it became clear towards the end of the second MIPR research program that information on the underlying physics of the various emitting devices (unavailable during either MIPR program) would be required in the future to yield reliable classifier performance. Nonetheless, it was demonstrated that even for a data set such as the new one, morphological measures provide performance comparable to or better than conventional techniques, and they do so with less demanding computational requirements.

The old data set consisted of 23 signals collected from six different emitters. The signals from one of the six classes are shown in Figure 4.1-1. It is important to note that the transients, although similar in appearance, are never in any way time-aligned or "registered". The classifier processing to be described is completely insensitive to the transient starting points within the extracted segments. (Similar insensitivity can be achieved by conventional means such as taking the magnitude of the transient's Fourier transform, but at a much greater computational cost).

A high-level block diagram of the classifier system is shown in Figure 4.1-2. As mentioned above, morphological and other advanced signal processing techniques were applied in the first two blocks. Variants of an efficient traditional classifier were used for the last block. For signal conditioning, two fundamental morphological filtering operators, *opening* and *closing*, were used to remove spikes and other undesirable temporal characteristics. Opening consists of a morphological *erosion* followed by a *dilation*. Closing consists of the concatenation of these two operators in the reverse order.

A two-dimensional example of erosion is shown in Figure 4.1-3. Like all basic morphological operators, erosion is based on a structuring *kernel* which is analogous to a finite impulse response (FIR) filter. The kernel is "slid" across an input signal or image just as is done with an FIR filter during convolution, except that each output sample is obtained from the minimum or maximum (for erosion and dilation, respectively) of the input samples under the sliding kernel, rather than the weighted sum of those samples (as in convolution).




Fig. 4.1-1. Typical raw signals (plotted with amplitude offsets).

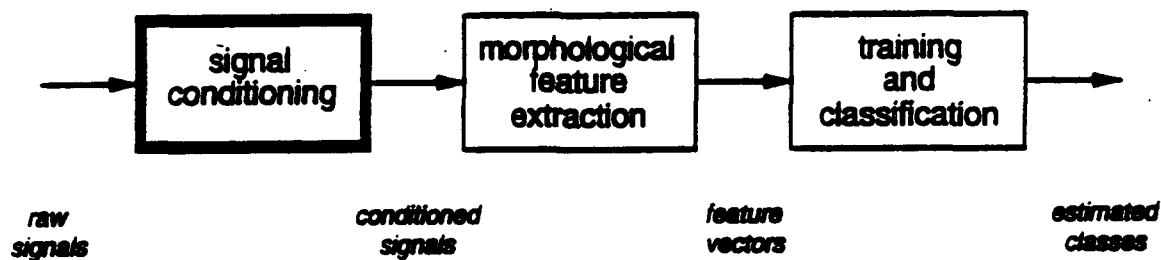


Fig. 4.1-2. Classifier system block diagram.


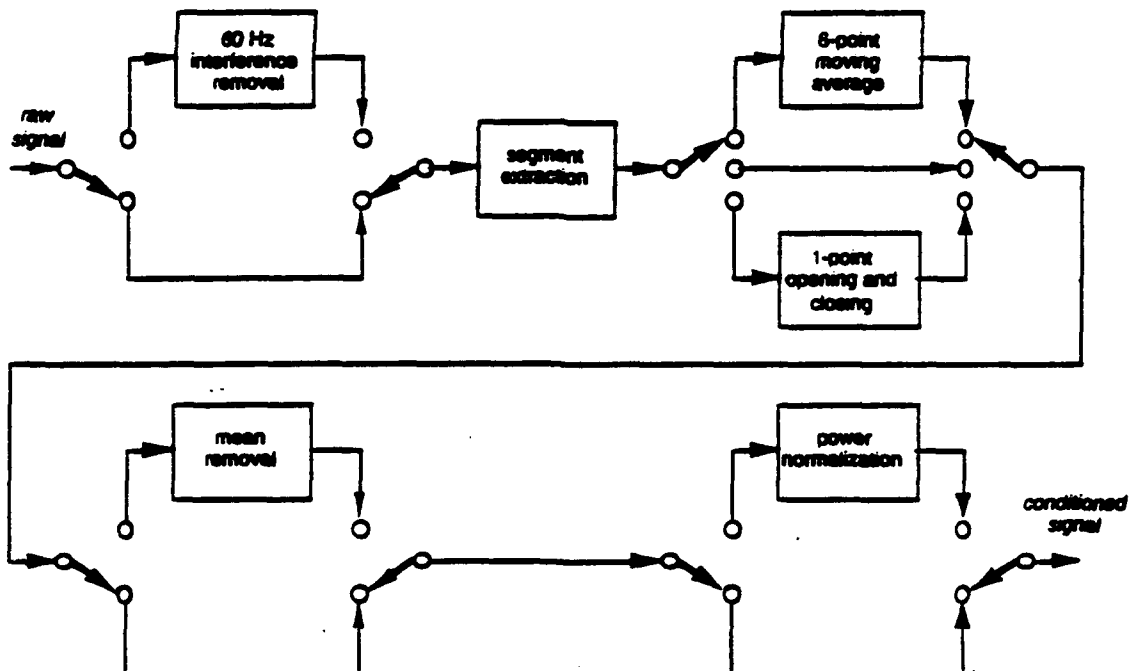


Fig. 4.1-3. Erosion process.



switches shown in optimal positions

Fig. 4.1-4. Signal conditioning experiments.


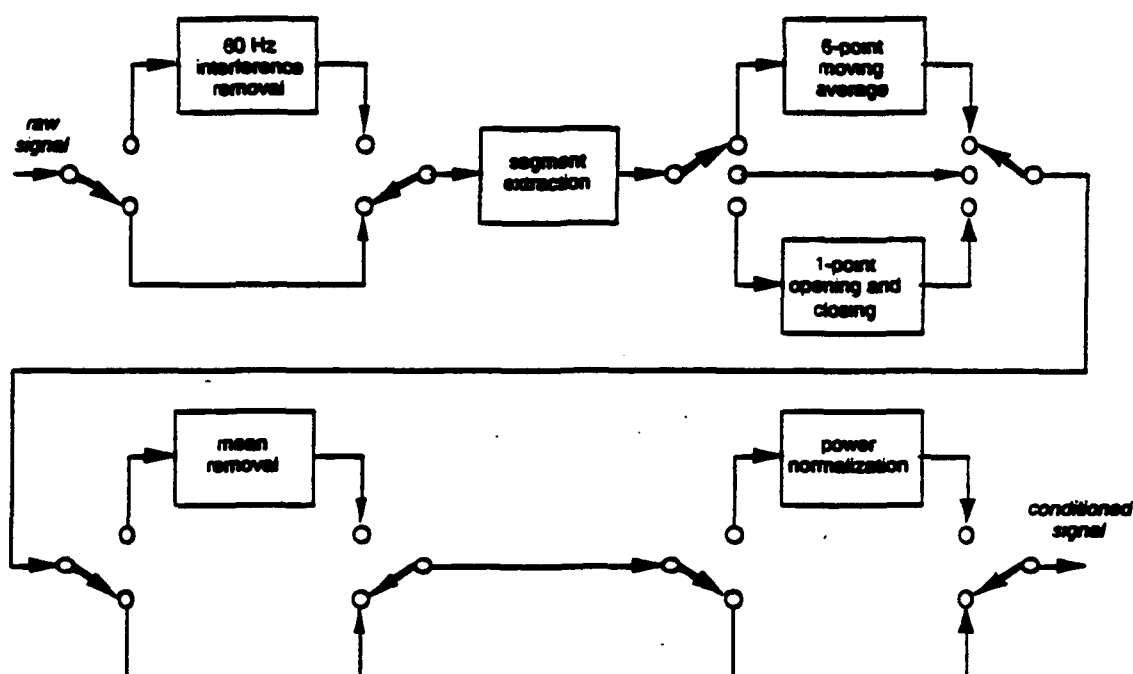


Fig. 4.1-3. Erosion process.



switches shown in optimal positions

Fig. 4.1-4. Signal conditioning experiments.

Many combinations of morphological and more traditional signal processing techniques were applied to the raw signals prior to feature extraction, as suggested in Figure 4.1-4. References to number of points indicate FIR filter or morphological kernel sizes. As mentioned earlier, the segment extraction was performed manually. The switches in the block diagram are shown in the positions which ultimately yielded the best classification performance. They show that the best signal conditioning was achieved using a 6-point moving average filter.




Fig. 4.1-5a. Effect of signal conditioning; before conditioning.




Fig. 4.1-5b. Effect of signal conditioning; after conditioning.

Class F:




Fig. 4.1-5c. Effect of signal conditioning; as input to feature extractor.

The effect of the preferred signal conditioning method is shown in Figure 4.1-5. The first two plots show the extracted raw signals for a single class before and after conditioning. Note that the various signals in these two plots have been time-registered and have had their means removed for visual comparison only. The third plot shows the conditioned signals as actually input to the feature extractors.

Given that linear filtering was judged to be best for signal conditioning, the dominant role of morphological processing was in feature extraction. For this purpose, variants of a higher-level morphological product known as *pattern spectra* were used. Like openings, closings, and other morphological operators, pattern spectra can be applied to data with any number of dimensions (in our case, just one). They are based on sequences of increasingly larger structuring kernels which are matched to expected waveform features. These kernels are successively applied to a given signal and the change in area under the signal stored as a pattern spectrum "bin" value (see Figure 4.1-6). By convention, the change in area due to larger and larger openings are plotted as positive values on the right side of a two-sided pattern spectrum. Negative values are plotted on the left for the corresponding closings. (Figure 4.1-6 only shows the openings spectrum.).





Fig. 4.1-6. Calculation of pattern spectrum.

Pattern spectra have a number of attributes which are relevant to signal source identification. They are sensitive to small signal substructure such as ringing and overshoot, and are unhampered by repetitive patterns in the data. They are also insensitive to macroscopic attributes such as temporal registration and overall signal amplitude. Like most morphology-based products, pattern spectra can be easily and rapidly computed in real-time using special-purpose hardware based on parallel architectures, and do not suffer from the dynamic range growth associated with most conventional techniques.



switches shown in optimal positions

Fig. 4.1-7. Feature extraction experiments.

Figure 4.1-7 summarizes the various feature extraction methods which were applied to the "old" data set. (A few additional techniques were applied to the new data set, as discussed later.) A new type of pattern spectrum was developed during the MIPR program and found to be the best method, among all those tried, for characterizing emitter signatures. It employed logarithmically spaced pattern spectrum (LSPS) bins for which kernel size was increased exponentially (rather than linearly, as done in the fine-grain (FGPS) and coarse-grain (CGPS) pattern spectra). The log-spaced pattern spectra for the signals from Figure 4.1-1 are shown in Figure 4.1-8.

Once candidate features were extracted from all test signals, they were fed into a traditional maximum-likelihood classifier. Such classifiers assume multi-variate Gaussian feature distributions characterized by hyper-elliptical contours of constant probability for each class i , as shown in Figure 4.1-9. In general, each estimate of the class ω_i to which a given feature vector most likely corresponds is based on the value of i which maximizes the log-likelihood function

$$\ln p(\omega_i) - \frac{1}{2} \left| \sum_i \left| -\frac{1}{2}(\bar{x} - \bar{\eta}_i)^T \sum_i^{-1} (\bar{x} - \bar{\eta}_i) \right. \right|. \quad (4.1.1)$$




Fig. 4.1-8. Example of log-spaced pattern spectra.


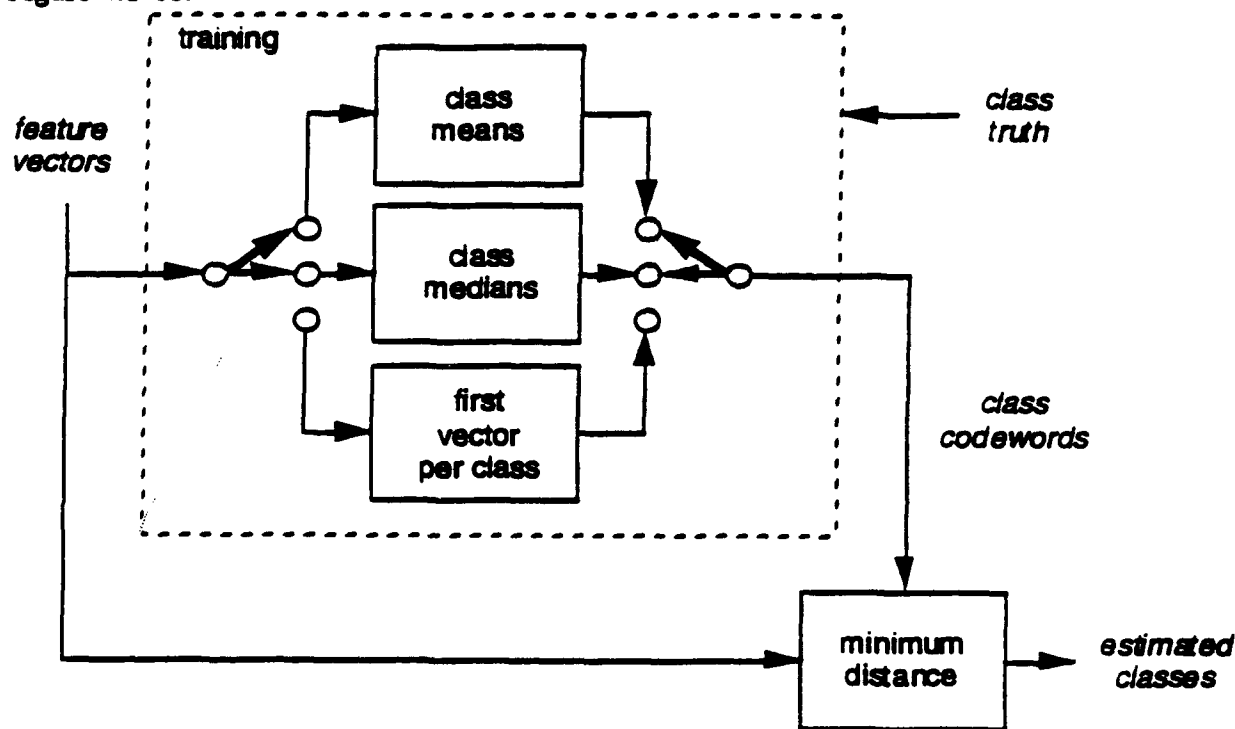


Fig. 4.1-9. Maximum-likelihood classification.

For simplicity, we assumed that the *a priori* class probabilities $p(\omega_i)$ and the covariance matrix determinants $|\Sigma_i|$ were equal, so that maximizing the log-likelihood function from above reduces to minimizing the normalized-axes distance function

$$(\bar{x} - \bar{\eta}_i)^T \Sigma_i^{-1} (\bar{x} - \bar{\eta}_i) \quad (4.1.2)$$

The minimum distance (MINDIS) classifier was trained in three different ways, as depicted in Figure 4.1-10.



switches shown in optimal positions

Fig. 4.1-10. Emitter classification experiments.

Mean and median feature vectors were found to be equally effective in characterizing specific emitter classes. Means are somewhat more attractive from a computational standpoint since they do not require sorting. Simply taking the first feature vector from each class's training set, although computationally inexpensive, was (not surprisingly) less effective. The results of apply the mean- (or median-) and first-vector training methods to eight different types of feature vectors are summarized in Tables 4.1-1(a) and 4.1-1(b). Only 14 features were necessary to achieve the best performance of 96% correct (highlighted by a box in the table) using the mean-based classifier. The other boxes in the table highlight signals which were misclassified.

	bs_1be_a	bs_2be_a	bs_3be_a	saw1saw	mm1cs_a	mm2cs_a	mm1mmmb	mm2mmne	mm3mmne	mm1mmnc	mm1mmne	mm2mmmb	mm2mmnc	mm3mma	mm3mmnc	mm1mmmb	mm1mmnd	mm1mmf	mm1mmg	mm2mma	mm2mmnd	mm3mmmb	mm3mmnd	% correct
TRUE CLASS:	a	a	a	b	c	c	d	d	d	e	e	e	e	e	e	f	f	f	f	f	f	f	f	
ESTIMATED CLASS BASED ON:																								
1. LSPS (22 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	
2. FGPS (20 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	
3. CGPS (20 features)	b	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	
5. 1 & 2 (42 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	
6. 1 & 3 (42 features)	b	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	
4. 2 & 3 (40 features)	b	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	
7. 1 & 2 & 3 (62 features)	b	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	


(a)

	bs_1be_a	bs_2be_a	bs_3be_a	saw1saw	mm1cs_a	mm2cs_a	mm1mmmb	mm2mmne	mm3mmne	mm1mmnc	mm1mmne	mm2mmmb	mm2mmnc	mm3mma	mm3mmnc	mm1mmmb	mm1mmnd	mm1mmf	mm1mmg	mm2mma	mm2mmnd	mm3mmmb	mm3mmnd	% correct
TRUE CLASS:	a	a	a	b	c	c	d	d	d	e	e	e	e	e	e	f	f	f	f	f	f	f	f	
ESTIMATED CLASS BASED ON:																								
1. LSPS (22 features)	a	b	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	91
2. FGPS (20 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	96
3. CGPS (20 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	96
5. 1 & 2 (42 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	96
6. 1 & 3 (42 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	96
4. 2 & 3 (40 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	96
7. 1 & 2 & 3 (62 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	96
8. 1 pruned (14 features)	a	b	a	b	c	c	d	d	a	e	e	e	e	e	e	f	f	f	f	f	f	f	f	91

(b)


Table 4.1-1 - MINDIS classifier results for old data set. (a) Using mean or median feature vectors for each class. (b) Using first feature vector for each class.

Versus signal from true class (mmm1mmmh, from class D):



(a)

Versus signal from estimated class (bs_1bs_a, from class A):



(b)

Fig. 4.1-11. Sole misclassified signal (mmm3mmme). (a) Versus typical signal from true class. (b) Versus typical signal from estimated class.

The misclassified 4% of the old data set comprised a single signal from among all 23. Figure 4.1-11 provides a visual comparison of this signal ("mmm3mmme", from Class D) with typical signals from its true class ("mmm1mmmh", also from Class D) and the estimated class ("bs_lbs_a", from Class A). The "time-constants" and other waveform structure associated with "mmm3mmme" are clearly much more similar to those in "bs_lbs_a" than those in "mmm1mmmh".

	bs_1be_a	bs_2be_a	bs_3be_a	saw1sawa	mnm1cs_a	mnm1cs_a	mnm1mnmnb	mnm2mnmne	mnm3mnmne	mnm1mnmnc	mnm1mnmne	mnm2mnmnb	mnm2mnmnc	mnm3mnmna	mnm3mnmnc	mnm1mnmnb	mnm1mnmnd	mnm1mnmnf	mnm1mnmng	mnm2mnmna	mnm2mnmnd	mnm3mnmnb	mnm3mnmnd	ggg1aa1	ggg1aa2	ggg1aba1	ggg1afa1	hhh1cla1	lll1lga1	iii1ggc1	iii1gdc1	iii1gde1	iii1gbl1	jjj1hcb1	lll1hcc1	% correct
TRUE CLASS:	a	a	a	b	c	c	d	d	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	q	q	q	q	q	a	i	i	i	i	i	i	
ESTIMATED CLASS BASED ON:																																				
1. LSPS (22 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	g	g	g	g	h	h	i	i	i	j	j	97
2. FGPS (20 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	g	g	g	g	h	h	i	i	i	j	j	97
3. CGPS (20 features)	b	a	a	b	c	c	d	d	a	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	g	g	g	g	n	h	i	i	j	j	94	
5. 1 & 2 (42 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	g	g	g	g	h	h	i	i	j	j	97	
6. 1 & 3 (42 features)	b	a	a	b	c	c	d	d	a	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	g	g	g	g	h	h	i	i	j	j	94	
4. 2 & 3 (40 features)	b	a	a	b	c	c	d	d	a	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	g	g	g	g	h	h	i	i	j	j	94	
7. 1 & 2 & 3 (62 features)	b	a	a	b	c	c	d	d	a	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	g	g	g	g	h	h	i	i	j	j	94	
8. 1 pruned (14 features)	a	a	a	b	c	c	d	d	a	e	e	e	e	e	e	e	f	f	f	f	f	f	f	f	g	g	g	g	h	h	i	i	j	j	97	

Table 4.1-2. MINDIS classifier results for extended old data set.

Morphological processing is attractive for the emitter classification problem not only due to its capability for shape-oriented signal discrimination, but for its ease of implementation and computational efficiency as well. To evaluate the computational complexity of our techniques and compare them with more traditional methods, let us define the following:

- the signal length, $L = 71$ to 501 (typically approx. 100)
- the number of bit per sample, $B = 16$
- the moving average window length, $M = 6$
- the maximum log-spaced pattern spectrum kernel size, $K = 64$
- the number of classes, $Q = 6$

Let us assume that an addition, subtraction, minimum, or maximum computation has a relative computational cost of one operation, while a multiplication requires B operations (essentially B additions). The unnormalized short-term moving average used for signal conditioning requires

$$C_{ms} = (M-1)L \text{ adds} \approx ML \text{ ops} \quad (4.1.3)$$

For each kernel size, k , the log-spaced pattern spectrum needs

$$C_{ps,k} = (k-1)L \text{ mins} + (k-1)L \text{ maxs} + 2L \text{ subs} + 2(L-1) \text{ adds} = 2kL \text{ ops.} \quad (4.1.4)$$

Thus, the total computation for the log-spaced pattern spectrum is

$$C_{ps} = \sum_{k=1,2,4,\dots,K} C_{ps,k} = \sum_{i=0}^{\log_2(K)} 2(2^i L) = 2(2^{\log_2(K)+1} - 1)L = 4KL \text{ ops} \quad (4.1.5)$$

The MINDIS classifier requires

$$\begin{aligned} C_d &= 2(\log_2(K)+1)Q \text{ muls} + 2((\log_2(K)+1)-1)Q \text{ adds} + (Q-1) \text{ mins} \\ &= 2QB\log_2(K) \text{ ops} \end{aligned} \quad (4.1.6)$$

The total computational cost is thus approximately

$$C = (M + 4K)L + 2QB\log_2(K) \text{ ops / signal} \quad (4.1.7)$$

For the values enumerated earlier, the total computation is dominated by feature extraction, i.e.,

$$C = 4KL \text{ ops} = 27 \text{ kops / signal} \quad (4.1.8)$$

For comparison, the computation associated with Fast Fourier Transform (FFT) based feature extraction is

$$\begin{aligned} C_{FFT} &= L\log_2(L) \text{ cplx - mul - adds} = 4L\log_2(L) \text{ mul - adds} \\ &= 4BL\log_2(L) \text{ ops} \end{aligned} \quad (4.1.9)$$

For the values above,

$$C_{FFT} = 43 \text{ kops / signal} \quad (4.1.10)$$

which compares favorably with the 27 kops/signal required for the morphological approach.

In summary, the performance of the linear filtering signal conditioning, morphological feature extraction, and minimum-distance classification worked well on the old data set. The preferred system correctly classified 22 (96%) of the 23 signals in the old data set. Comparable results (97% correct) were achieved after also including 12 new sponsor-provided signals. All processing elements, comprising unnormalized short-term moving averages for signal conditioning, log-spaced pattern spectra for feature extraction, and the MINDIS classifier based

on class-wise vector means, were computationally inexpensive. The morphological transform typically reduced signature data by 90% and obviated the need for time registration.

Given the success above, similar techniques were applied to the new data set. Unfortunately, the class "truth" provided with the new data appeared questionable based on inspection of the plotted waveforms: signal structure varied widely between signals in each class and, more importantly, there were many suspicious similarities between signals in *different* classes. We therefore decided to reassign the data to a new set of classes using sophisticated, well-known classifier tools employing neural network analysis, and then apply our classifier system to the data with the new, neural net truth. In addition to assigning the signals to new classes—discarding sponsor-defined class assignments—the neural network program was intended to determine the true number of classes represented by the new data set. Once the classes were reassigned, the preferred preprocessing and classification algorithms which were successful on old data were applied to the new data, and alternate feature extraction methods were developed and tested.

Neural network analyses are categorized by two basic paradigms, supervised and unsupervised. Supervised neural nets learn and generalize existing classes and are most suitable when the class exemplars are well known. In sharp contrast, unsupervised neural nets are used to discover classes in unclassified data and thus are suitable when the underlying classes are undetermined or vague, as was the case for the new data set.

The neural network program which we selected for use as our "higher authority" for class reassignment was SOM-PAK (the "Self-Organizing Map Program Package"), one of the most readily available unsupervised clustering networks. It was built by the team which developed the theory [KOH0] at the Helsinki University of Technology Laboratory of Computer and Information Science. Its one disadvantage is that it does not automatically estimate the number of classes in the data and thus requires a hypothesized number of classes as an input.

The process by which we reassigned classes using SOM-PAK and evaluated our own classifiers is summarized in Fig. 4.1-12. As evident in the figure, only log-spaced pattern spectra were input the neural net classifier, whereas all types of feature vectors were used for testing our classifier.

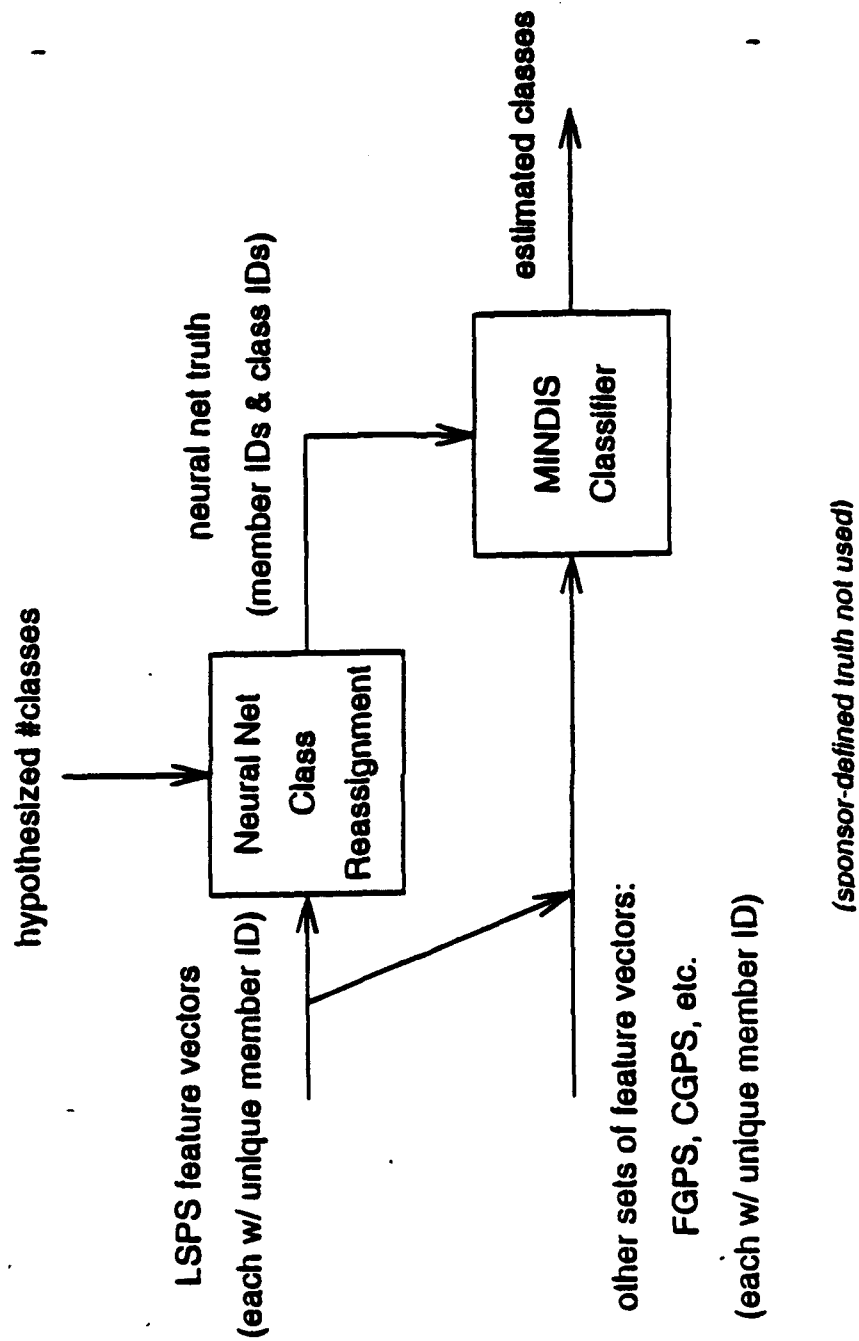


Fig. 4.1-12. Classifier performance evaluation using class reassignment.

The preliminary experiment conducted using SOM-PAK was to search for any degree of correlation between the sponsor-defined and neural net class assignments for the new data set. It was hypothesized that the number of classes was 20, a number somewhat larger than the number of sponsor-defined classes (11), with the assumption that excess classes would be sparsely populated. Table 4.1-3 consists of tallies of sponsor-defined class members in each neural-net-defined class. There is no apparent correspondence between the two sets of class assignments. Had there been any, there would have been only one relatively large value in each row (the first row contains an 11, a 6, and a 5; the second contains a 13 and a 9; etc.).

		sponsor-defined class											TOTAL
		a	b	c	d	e	f	g	h	i	j	k	
neural net class	1	1	3	2	11	3	3	5	0	2	6	2	38
	20	0	1	4	0	0	9	13	2	0	0	2	31
	12	3	0	0	0	1	1	0	0	0	1	4	10
	7	1	1	0	2	2	1	1	0	1	1	0	10
	5	1	1	0	1	1	2	0	0	3	0	0	9
	4	2	2	0	1	1	1	0	0	0	0	2	9
...	

Table 4.1-3. Tallies of sponsor-defined class members in each neural-net-defined class.

The next preliminary experiment was to estimate the number of classes spanning the new data set manually, with the aid of SOM-PAK. Table 4.1-4 illustrates the distribution of the signal population over neural net classes for varying hypothesized numbers of classes. For comparison, boxes are shown around values from Table 4.1-3. In viewing the number of members per class, we are looking for a cutoff (a row number) below which the remaining classes are lightly populated. There appear to be two dominant (heavily populated) classes, and about five classes total plus individual, random (unassociable) cases.

In the absence of clear information on the number of underlying classes in the new data set, we decided to evaluate classifier performance for various numbers of classes, as shown in Table 4.1-5. The values shown are percent correct using various types of normalized pattern spectra (log-spaced, etc.) with training based on mean feature vector codewords. Classes were reassigned by running the neural net on the LSPS data only once per hypothesized number of classes. The best performance (76% correct) was obtained by concatenating the LSPS and FGPS feature vectors, for a total of 42 features.

class	hypothesized number of classes									
	12	14	16	18	20	22	24	26	28	30
1	62	56	52	46	38	33	33	30	28	23
2	38	35	35	34	31	27	25	23	23	22
3	15	14	12	10	10	11	12	10	10	11
4	11	10	12	10	10	10	10	10	10	9
5	11	10	11	9	9	10	10	9	10	9
6	11	9	11	9	9	9	9	8	7	8
7	9	9	7	9	8	9	7	8	7	8
8	8	8	7	8	8	8	7	8	7	7
9	6	7	7	7	8	7	7	8	7	7
10	5	7	6	7	8	7	6	7	7	7
11	2	7	6	7	7	7	6	7	6	6
12	2	3	5	6	6	6	6	6	6	6
13	0	3	4	6	6	6	5	5	6	6
14	0	2	4	4	5	4	4	5	5	5
...					...					

Table 4.1-4. Distribution of population over neural net classes.

class	hypothesized number of classes									
	12	14	16	18	20	22	24	26	28	30
1	62	56	52	46	38	33	33	30	28	23
2	38	35	35	34	31	27	25	23	23	22
3	15	14	12	10	10	11	12	10	10	11
4	11	10	12	10	10	10	10	10	10	9
5	11	10	11	9	9	10	10	9	10	9
6	11	9	11	9	9	9	9	8	7	8
7	9	9	7	9	8	9	7	8	7	8
8	8	8	7	8	8	8	7	8	7	7
9	6	7	7	7	8	7	7	8	7	7
10	5	7	6	7	8	7	6	7	7	7
11	2	7	6	7	7	7	6	7	6	6
12	2	3	5	6	6	6	6	6	6	6
13	0	3	4	6	6	6	5	5	6	6
14	0	2	4	4	5	4	4	5	5	5
...					...					

Table 4.1-5. MINDIS classifier performance relative to neural net class assignments.

<i>features</i>	<i>#features</i>	<i>%correct</i>
pruned LSPS	14	38
log-spaced pattern spectrum (LSPS)	22	38
CGPS of complex envelope	20	38
fine-grain linear-spaced pattern spectrum (FGPS)	20	37
fast Fourier transform (FFT) magnitude	257	32
coarse-grain linear-spaced pattern spectrum (CGPS)	20	31
FGPS of complex envelope	20	29
LSPS of complex envelope	22	26
FGPS of FFT magnitude	20	25
LSPS of FFT magnitude	22	23
OGPS of FFT magnitude	20	22
fractal dimension	4	19
linear predictive coding (LPC) modeling error	10	12
LPC modeling error of complex envelope	10	11
fractal dimension of complex envelope	4	10

Table 4.1-6. Exploration of alternate feature extraction methods.

Alternate feature extraction methods were also investigated. The original (i.e., sponsor-provided) truth was used for this purpose to provide a fair comparison between the LSPS, which had also been used for class reassignment, and other types of features. Table 6 compares the size and effectiveness of the various feature vectors. All pattern spectra vectors were unnormalized. As with the old data set, the best performance (38% correct for the new data set) was achieved with the pruned LSPS comprising only 14 features. Other, much more computationally intensive signature measures such as FFTs and parametric (LPC) modeling did not improve performance.

4.1.3 Year 3 Effort (Greenbelt Facility)

4.1.3.1 Experiments

In year 3, the second database of emitter data was processed with the Gabor software. This data consisted of 1944 files which contained several instances of data per file. Many GSPS runs were made using signal data and other functions as windows. In these cases, the minimum dimension resulting from the Gabor processing was higher than in previous runs, because the second database contained much more complicated signals. The identification of which portion of each file was the signal of interest was difficult, if not impossible without knowledge of the source of the data.

4.1.3.2 Results and Observations

After many attempts at classifying elements of the second database, it was determined that the second set of data was unreliable, and the "ground truth" could not be known due to errors made by another contractor when generating the data.

4.1.5 Future Work

It is clear that future work on the specific emitter identification problem will be most effective if the physics of the emitting devices are analyzed and quantified. This was not possible during the programs discussed above due to the limited availability of supplemental data. It will probably be necessary to characterize the transmission channel through which the signals were obtained, and perhaps the data acquisition process, as well.

Armed with all of this information, it should be possible to implement various rapid signal classification processing techniques. Work performed using the Gabor expansion signal dimension was promising, but was not further testable under this year's effort. Extensions might combine Gabor methods with those based on pattern spectrum and other morphological discriminants. The excellent results (96% correct) obtained for the old data set of 23 signals suggest that the log-spaced pattern spectra will provide an effective, efficient method for feature extraction, in terms of both classification reliability and feature vector compactness. Other morphology-based features should also be investigated since they can be tuned to signal "appearance", *i.e.*, they can discriminate various shapes in transients in a manner similar to what human analysts do. Low cost, real-time implementations are also possible, and should be developed and analyzed with respect to size, cost, and power requirements.

4.2 AUTOMATIC TARGET RECOGNITION

4.2.1 Problem Definition

4.2.1.1 Statement and Importance of the Problem

The detection of small targets in clutter is a problem of critical importance to wide-area, long-range, InfraRed Search and Track (IRST) surveillance applications. In these applications, the targets (aircraft and missiles) are typically unresolved and therefore appear in sky, sea and terrain backgrounds in only a few resolution cells. A resolution cell is the partitioning of the search space in azimuth and elevation. The number of resolution cells that the target spans is predominantly determined by the system response of the sensor system.

Figure 4.2-1 illustrates a typical IR system. The target and background comprise a scene. The optical system images the scene onto the detector array which in turn samples the intensities of the image. The detector outputs are then digitized and stored in a memory array as the digitized representation of the image. The optical system is characterized by its point spread function (PSF). This function is the response of the optics to unit 2-D impulse functions or point source inputs. Circular aperture diffraction limited optical systems typically have an impulse response with the intensity in the focal plane in a radial direction from center given by:

$$E(r) = E_0 [2J_1(m) / m]^2. \quad (4.2.1)$$

where:

E_0 = peak illuminance


$J_1(m)$ = first order Bessel function

$m = 2\pi (NA) r / \lambda$

$NA = (2 F\#)^{-1}$ is the numerical aperture of the system

The system aperture acts as a spatial filter with the above characteristic, bandlimiting the spatial frequencies passed by the optics. In essence, the pupil function that describes the aperture is convolved with the scene input to form an image on the detector array. The result of the spatial filtering operation is a reproduction of the scene with the spatial spectral content shaped by the optics. In well designed IR systems, the detector array makes use of the lowpass

characteristics of the optics and samples the image intensities at a spatial rate of at least two times the cutoff frequency of the optics. This is done to prevent aliasing artifacts. The result is an array of data representing the image intensities as sampled by the detector array. This is the data that must be analyzed to find targets.



The above IR system showed a two-dimensional focal plane detector array to simply illustrate the concept. For several reasons including detector technology, upgrade potential and cost, many IR systems employ linear detector arrays and a scanning mechanism which scans the image across the array to provide the second dimension to the data. This is the mechanism used in the IRAMMP sensor described below. The angular size of a single detector in image space usually defines the instantaneous field of view (IFOV) of the sensor system and the angular extent of the image scanned defines the total field of view (TFOV). The amount of time that is required to cover the search area with the IR sensor and process the data is known as a scan period.

Because of the nature of the threats, it is sometimes necessary to declare a target within a single scan period. Thus, this application not only requires an efficient processing approach but also a robust detection performance in the presence of clutter using only single scans of data. HorizonIRST systems may not have to contend with exceedingly strong cloud and land clutter or water glint. This of course is a function of the extent of elevation coverage the sensor provides about the horizon. In most cases, one may expect low contrast situations in the vicinity of the horizon. In addition, IR clutter fields are in general both spatially and temporally nonstationary. This fact weighs heavily on any design of a best filter for detection processing.

The background data that has been used to test filter algorithms has been extracted from the Navy's Infrared Analysis Modeling and Measurements Program (IRAMMP) database. The IRAMMP database contains an extensive collection of radiometric data including a variety of backgrounds. The IRAMMP sensor used to gather this data is a dual-band radiometer which records IR backgrounds in the 3-5 and 8-12 micron bands simultaneously. For each band, a linear staggered contiguous focal plane array consisting of 120 detector elements is used. The elements are staggered to allow alignment of the detector active area edges in the vertical direction. This allows for the fact that if the detectors were mounted in a vertical column with their edges almost touching, there would be a "dead space" between them where no intensity would be measured.

The detectors are electrically sampled in synchronism with the scanning mechanism so that successive samples of the image taken from the staggered elements are registered to the same horizontal location in the image. The instantaneous field of view (IFOV) of each detector element is 0.23 milliradians vertically by 0.22 milliradians horizontally. The impulse response of the optical system matches this closely. The optics were measured to have 85% of the point source throughput within 1.09 IFOV for longwave and 1.18 IFOV for midwave. The total field of view is 1.6 degrees vertically by 5.6 degrees horizontally. The horizontal scan is electrically sampled 1480 times resulting in a spatial angular spacing of 0.0664 milliradians per sample. This corresponds to a 3.3 times horizontal oversampling of the image by each detector element. The sampling in the vertical direction is not as good, being only 1 times (approximately). The result is vertical undersampling of the image, with the potential for spatial aliasing artifacts to occur in the vertical dimension. This is not a problem for algorithms that work only on a single scan line basis, but the aliasing artifacts could cause a problem for 2-dimensional algorithms. The vertical problem could be alleviated by detector arrays staggered to give at least a 50% overlap of the image as it is scanned across the arrays.

The measured IRAMMP detector parameters are as follows: The noise equivalent temperature (NET) difference is 0.047 °C for the mid-wave and 0.032 °C for the long-wave. The noise equivalent irradiance (NEI) is 2.6×10^{-14} W/cm² for mid-wave and 2.6×10^{-13} W/cm² for long-wave. The dynamic range of the data is 84 dB which corresponds to 15 bits. The uniformity of response across the detector elements was not measured. It is known that detector point response across the element surface can vary quite a bit from edge to edge. Since the optics were well matched to the detector size this should not cause any problems. Inaccurate data results when the optics create a spot size much smaller than the element size.

The IRAAMP background data can have synthesized or real target data overlaid on it digitally for testing various filter processes. Several different types of detection filters have been used to date with varying degrees of success. For the detection of weak targets in homogeneous IR backgrounds, optimal, linear, finite impulse response (FIR) filters have been developed. One such filter is the least-mean-squared filter. The nonstationarity of IR clutter backgrounds however, results in degraded performance for linear filters. To overcome this deficiency, nonlinear techniques for point target detection have been proposed. Within the past several years, a class of nonlinear signal processing algorithms, collectively known as morphological filters, has been applied to machine vision problems and, more recently, to target detection problems. These algorithms respond to the size, shape and orientation of imaged objects using computationally efficient logic. The result is an improvement in real-time response and target detection capability over linear filter techniques.

Recent work has shown that a filtering process based upon a matrix formulation of the Gabor transform holds promise as a mechanism for nonlinear processing for the IRST problem. The Gabor transform can be used as a pre-filter for the morphological processing. The resulting combination yields a process that is very much higher in performance than linear algorithms used in the IRST application. The next section describes the formulation of the algorithm used and details the results of experiments with the Gabor transform algorithm.

4.2.1.2 Current Techniques

4.2.1.3 Applicability of Gabor Transform

The Gabor transform is suggested for the ATR application because it inherently resolves the data subjected to it into time- and frequency shifted replicas of a window function that may, within limits, be chosen arbitrarily. When the data is presumed to include replicas of signals, a Gabor transform with a high degree of time resolution may be appropriate as an analysis mechanism. For unresolved target IRST, the choice of a window is simplified by the fact that the system response to any one target is simply the point spread function of the receive optics.

The optics of the IRST system define, in general, the shape of the response that is expected at the receiving end of the system. Due to the fact that the system oversamples nine times for each azimuthal rotation of one degree, the shape of the received response to a single point target will be well modeled by a Gaussian pulse as shown in Figure 4.2-1.

4.2.2 Experimental Results

4.2.2.1 Definition of the Test Problem

We will now proceed to the definition of the problem that we used for our Gabor transform experiments. Fig. 4.2-1 is a cloud scene from an IRST return upon which we superimposed a stencil of targets depicted in Fig. 4.2-2. The resulting image is given by Fig. 4.2-3, in which we notice that the targets are not discernible to the naked eye. This is due to the high dynamic range associated with this particular type of scenario. Next, we extracted one line in elevation that contains about 20 targets buried in the cloud cover, and used it as a waveform (Fig. 4.2-4). The data was then put through a 21-tap lowpass filter and the mean was subtracted out, resulting in the waveform of Fig. 4.2-5, which became the starting signal for our 1-dimensional Gabor experiments. The question that can now be addressed is whether or not the two-dimensional time-frequency representation that results from the Gabor transform aid us in localizing the targets. Preliminary experimental results and detailed description of the experiments are discussed in the following sections.




Fig. 4.2-1. Cloud cover section.




Fig. 4.2-2. Targets.



Fig. 4.2-3. Clouds plus targets.

4.2.2.2 Experiment Description

After importing the 1024 data point signal described above into the GSPS workstation, we proceeded to test different windows and time-frequency splittings that might yield results deemed promising in either signal cleanup or target arrival detection. To this end we began by choosing the time-frequency split to be 256 time points and 4 frequency points, and used the Longwave point response of Fig. 4.2-6 as our analysis window. In the two-dimensional map of Gabor coefficients, Fig. 4.2-7, we can clearly distinguish isolated amplitude peaks corresponding to some of the target locations along the off-DC frequency lines. One will notice that some of the spikes have smaller amplitudes than others and some of the targets are not really well

represented due to time misalignment of the window with respect to the signal. Since we have chosen the number of frequency points to be 4, in the critically sampled case it will be sufficient to perform three successive single-point shifts of the signal to guarantee that in at least one of the time shifts is as well aligned as possible with the chosen window. Figs. 4.2-8 through 4.2-10 represent the three time shifts, in which we can clearly see that different targets line up better than others for a given shift. Notice that targets that initially lined up best with the response window showed poorly in the shifts. Each signal of interest showed up most clearly in a certain time shift.




Fig. 4.2-4. Unfiltered signal.




Fig. 4.2-5. Filtered signal containing the targets.




Fig. 4.2-6. Longwave response used as a window.




Fig. 4.2-7. Gabor coefficients at 0 shifts.




Fig. 4.2-8. Gabor coefficients at 1 shifts




Fig. 4.2-9. Gabor coefficients at 2 shifts.




Fig. 4.2-10. Gabor coefficients at 3 shifts.

In order to address the automatic target localization problem, we decided to perform some simple nonlinear processing to see if we could isolate the coefficients that correspond to the target locations. This procedure was done in the following manner. First we decided to use the power of MATLAB to aid us in this task, so we saved the four sets of coefficients from GSPS in MATLAB format and exported them into that application to perform our nonlinear processing. Once we had the coefficients we processed the coefficients according to the following algorithm:

Let a_0 , a_1 , a_2 , a_3 be the set of Gabor coefficients corresponding to 0, 1, 2, and 3 time shifts of the signal, respectively. Then the following pseudo-code describe the process of combining Gabor coefficient values across the shifts:

```

for i=1 to 256
  for j=1 to 4
    sum(j,i)=|a0(j,i)|+|a1(j,i)|+|a2(j,i)|+|a3(j,i)|
  end
end
maxval = max ( sum(j,i)) for all i,j
level = .05 * maxval
for i = 1 to 256

```

```

if any ( |a0(2,i)|, |a1(2,i)|, |a2(2,i)|, |a3(2,i)| > level )
    proc_coef(j,i) = a0(j,i) for all j
end
end
plot(proc_coef).

```

After the processing has been done, import the array proc_coef of processed coefficients back into the GPS system and reconstruct. The coefficient map is shown in Fig. 4.2-11 and the reconstructed signal is shown in Fig. 4.2-12. Notice that all except one of the targets were isolated, and most of the clutter that was originally in the signal has been removed.




Fig. 4.2-11. Processed Gabor coefficients.




Fig. 4.2-12. Signal reconstructed from the nonlinearly processed coefficients.

Clearly it will take some more sophisticated processing to isolate all the targets and decrease the false alarms, but as a proof of concept this demonstrates that the Gabor transform together with some nonlinear processing techniques can be of help in the localization of targets in clutter. Of course more work will have to be done to either theoretically or experimentally find the optimal window as well as the optimal grid distribution and signal length in such a way that the maximal number of targets is detected and the minimal number of false alarms is recorded. Other ways to proceed might be to use the oversampled Gabor representation to obtain the Gabor Coefficients, preprocess the signal further and zero out all the points that are negative, and to use the optimization methods described in section 4.3 to find an optimal window that is some kind of average of the two window responses that can be expected.

4.2.2.3 Other Experiments

To ascertain that the above observations were not artifacts of the uniform spacing, or results that were somehow inherent to the cloud structure, we also ran the following experiments using the same analysis window, the same time-frequency grid distribution, and the same four time shifts in order to exhaust the time alignment problem. We took the same cloud sample as before but without any targets in it, and used the same analysis window to obtain Gabor coefficients for zero to three time shifts. The unfiltered and filtered cloud signals are given in Figs. 4.2-13 and 4.2-14, while the time-frequency maps are given by Figures 4.2-15–4.2-18. As can be clearly seen, the structure that is present in the off-DC lines when the targets are present is no longer there when the targets are missing, *i.e.*, the algorithm does not generate false alarms.




Fig. 4.2-13. Cloud unfiltered.




Fig. 4.2-14. Cloud filtered.




Fig. 4.2-15. Gabor coefficients at 0 shifts.




Fig. 4.2-16. Gabor coefficients at 1 shift.




Fig. 4.2-17. Gabor coefficients at 2 shifts.




Fig. 4.2-18. Gabor coefficients at 3 shifts.

The next experiment that was performed was the verification that the high amplitude Gabor Coefficients were not artifacts of the uniform spacing of the targets.




Fig. 4.2-19. Clouds with non-uniformly spaced targets.




Fig. 4.2-20. Gabor coefficients at 0 shifts.




Fig. 4.2-21. Gabor coefficients at 1 shift.




Fig. 4.2-22. Gabor coefficients at 2 shifts.




Fig. 4.2-23. Gabor coefficients at 3 shifts.

Even though the window succeeded in pulling out the targets, the coefficient map is not as sharp as the map from the uniformly spaced coefficients. We will now see in Figs. 4.2-25–4.2-28 that the use of the Mediumwave response window of Fig. 4.2-24 will do a much better job of pulling out the same targets.




Fig. 4.2-24. Mediumwave response used as a window.




Fig. 4.2-25. Gabor coefficients at 0 shifts: MW window.




Fig. 4.2-26. Gabor coefficients at 1 shift: MW window.




Fig. 4.2-27. Gabor coefficients at 2 shifts: MW window.




Fig. 4.2-28 Gabor coefficients at 3 shifts: MW window.

One can easily see from the new set of pictures that some of the Gabor coefficients that correspond to target locations have been magnified, and some of the clutter due to spreading has been eliminated by the use of the (slightly) different window. Between the two representations, all of the targets that are present in the signal appear as large coefficients in the off-DC frequency lines. This certainly gives a coefficient map that is very similar to the one obtained from the uniform target spacing, therefore, we conclude that the spiky zones along the off-DC lines are a true manifestation of the presence of the target and not some artifact due to target spacing.

We also tried some experiments based on the following idea. In the impulse responses resulting from the longwave, the target will always be brighter than the clutter (this is not necessarily the case with the midwave responses), and hence all targets will have at least a portion which exhibits positive values. This implies that we might be able to take the filtered signal, arbitrarily zero out the points which have negative values, and process as before with all the necessary time shifts. The result of this extra signal preprocessing step should be that some of the clutter will already be removed before applying the transform, thereby cleaning up the transform coefficient map and increasing the visibility of the signal. For all the experiments described below we used the Longwave response as the window.

The result of this zeroing is depicted in Fig. 4.2-29. We then chose a splitting of 2 frequency and 512 time points and obtained the Gabor coefficient map for both of the significant time shifts, the result of which is shown in Fig. 4.2-30 (a) and (b). Analysis using a grid of 4 frequency and 256 time points was also performed. The resulting Gabor coefficient map for the four relevant time positions is displayed in Fig. 4.2-31 (a) - (d). Comparison of the Gabor coefficient maps obtained from the preprocessed and unprocessed signals indicates that the latter tend to spread the coefficient mass more across frequency and consequently cause the target display to become more shift invariant.






Fig. 4.2-29. Signal resulting from the removal of the negative values of the original filtered signal.




(a)




(b)


Fig. 4.2-30. Gabor coefficients for the processed signal with $M = 512$ and $N = 2$. (a) Coefficients for the unshifted signal and (b) coefficients for the signal after time shifting once to the left.




(a)



(b)




(c)




(d)

Fig. 4.2-31. Gabor coefficients for the processed signal with $M = 256$ and $N = 4$. (a) - (d) correspond to the coefficients for the signal after time shifting 0 - 3 points to the left respectively.



(a)



(b)

Fig. 4.2-32. Oversampled Gabor coefficients for the processed signal with $M = 1024$ and $N = 2$. (a) Coefficients for the unshifted signal and (b) coefficients for the signal after time shifting once to the left.




Fig. 4.2-33. Oversampled Gabor coefficients for the processed unshifted signal with $M = 1024$ and $N = 4$.




Fig. 4.2-34. Oversampled Gabor coefficients for the processed unshifted signal with $M = 1024$ and $N = 8$.

4.2.2.4 Results and Conclusions

The experiments performed here indicate potential for Gabor analysis as a component of unresolved target detection for theIRST problem. Synthetic targets representing replicas of the receive optics point spread function were distributed both uniformly and nonuniformly across an observed cloud background, and were found to be visible in processed Gabor expansions that favored time resolution as opposed to frequency resolution. There is a robustness in the results that indicates stability of the findings with respect to the details of the approach.

4.2.3 Future Work

Extensions of the current work would emphasize several things, among them further theoretical studies to gain an explanation of the moderate successes observed in the reported experiments. Because of the time resolution inherent in the oversampled approach, further experiments with it are warranted also. Application of the techniques to additional data sets could be expected to enhance our understanding of the performance as well as uncover areas where the techniques shown here are either inadequate or in need of support from other methodologies. In particular, we are anxious to investigate algorithms which combine the power of Gabor methodologies in producing data that exhibits signal detection in a few large coefficients and morphological filtering to aggregate the observed points and replace some "human operator" functions with machine capability.

4.3 MINIMUM DIMENSION GABOR

Under contract to DARPA, AAEC carried out Phase I and II SBIR studies on the topic of "Minimum Dimension Gabor Representations." This work drives at obtaining Gabor expansions of a set of signals in which the majority of the information about the signals is concentrated in a few large expansion coefficients. Methods of nonlinear optimization theory are employed in this quest. The work carried out under these contracts produced very positive results leading to two conference publications [SWEE, ORR], but did not complete the needed research. Some of the research extensions were carried out under this contract, and their results are presented below.

In the following paragraphs we summarize the problem statement, the theoretical results produced under the Phase I SBIR and the numerical algorithms and results from Phase II. Following this the new work performed under this contract is described.

4.3.1 Summary of SBIR Contract Effort

4.3.1.1 Problem Description

Series expansions of signals in which significant features of the signal are captured in a few large coefficients are desirable. This work shows that given a collection of signals, it is possible to find Gabor representations for these that are maximally concentrated in time-frequency space. The problem addressed is: given a signal set, find the window function of the Gabor expansion that minimizes an "average dimension" of the signal representations relative to that window. The dimension measure employed is entropy based and related to the quantum-mechanical technique where one interprets expansion coefficients as probabilities.

An iterative algorithm based on partial derivatives of the signal set dimension with respect to the expansion function was used to evaluate effectiveness of several nonlinear optimization algorithms in finding an optimum window.

4.3.1.2 Phase I Accomplishments

The key results of Phase I are presented here, and the reader is referred to the final report and a conference publication for details [AAEC91(2), ORR]. Key findings include:

- Proof that the notion of "dimension" is, for a single signal, completely arbitrary, depending wholly upon the choice of representation and not at all on the signal;
- A convincing demonstration that where a *set* of signals is concerned, the structure of the set bears an inherent relation to a *dimension* that can be assigned—this dimension remains a

function of the representation, but can be minimized over quite general subsets of all basis functions;

- Recognition that the constrained structure of time-frequency basis sets such as the Gabor, affine wavelet, *etc.*, make these ideal classes of bases over which to carry out the minimization of set dimension;
- A determination of a reasonable set of requirements for the definitions of both signal and signal set dimensions;
- A derivation from these requirements of a *signal* dimension formula that is unique to within a single parameter and a *set* dimension formula that is then unique to within the assignment of a weighting function—the formula combines the quantum mechanical relationship of expansion coefficients to probabilities and the information theoretic notion of the entropy of a probability distribution;
- Successful implementation of the signal dimension measure within AAEC's Gabor processing software testbed (GSPS);
- A useful, though incomplete, classification of signal set *types* in a manner that motivates physically meaningful choices of weighting functions for dimension assignment;
- An extensive compilation of examples of basis set *constraints* that would allow optimization over more restricted subsets of Gabor or other representation families;
- Development—via use of *Poisson summation* and the ambiguity function—of new characterizations of Weyl-Heisenberg frames (the immediate generalization of Gabor representations) that should make it easier to perform dimension optimization under constraints within these structures.

These achievements set the starting point for Phase II. Having successfully formulated the dimension concept and the associated minimization problem, the immediate need was to discover the extent to which analytic machinery can be brought to bear on solving the difficult nonlinear optimization problem we have created. For example, in basis systems that are defined by a single window function, which include most of the cases of interest—Gabor, Weyl-Heisenberg, affine wavelets, *etc.*—the solution consists of determining an optimum window. If there are sufficiently restrictive constraints placed upon the class of eligible windows, it should be possible to obtain solutions that not only minimize the cost function, but have other desirable behaviors as well. If we can obtain analytic answers for some simple but nontrivial cases, the utility of the techniques will be enhanced. To conclude this aspect it is important to achieve a software capability that implements the optimization algorithms that are found.

4.3.1.3 Phase II Accomplishments

Phase II concentrated mainly on application of a concept put forth under the prior phase [AAEC92(2)]. New theoretical developments were limited to only those needed to carry out the optimizations or to interpret results. The main theoretical idea is of course the development of the discrete domain cost function for optimization and the choice of optimization variables. Carrying out the optimizations over the biorthogonal function has proven quite successful, and has eliminated an intermediate computational burden, one that could have been quite expensive, were the window itself directly optimized. We see nothing in the results to prompt any revision our underlying methodology.

The final theory topic is the evaluation of computational complexity of the dimension functions. We found that a dimension or gradient evaluation for one signal is an $O(P^2)$ process (recall that P is the number of points in a signal), whereas a Hessian computation is $O(P^3)$. These are reasonably high orders—think of the effort people will expend to replace a traditional order $O(P^2)$ Fourier transform with an FFT. No pressing need to rely on great numbers of Hessians has been uncovered, so we can probably say with reasonable confidence that the per-signal computation caps at $O(P^2)$. No analysis of the adaptive algorithms was attempted since there are so many variants, and the run time is data-dependent.

Initial experiments using a local gradient method were very successful in demonstrating that iterative techniques can converge to a good solution. A number of other things were observed from these as well. Convergence to a suboptimum point, *i.e.*, a local minimum, was seen occasionally, and in some circumstances the answer was a very plausible candidate. In some cases the converged biorthogonal caused the Gabor coefficients of the signal to be dominated by a single large coefficient located somewhere other than the origin; the resulting window was either incalculable or did not resemble the signal. This circumstance should be regarded as a successful intermediate result from which the search for a more properly behaved solution can be sought. Only partial success in carrying on from this point was achieved, leaving this circumstance as one of the candidates for further work.

The more appropriate experiments using the optimization techniques of Matlab and NAG yield further confirmation that the process would work, but some more surprising results were found. Most startling among these is the observed convergence to highly unstable answers. There exist cases for which a solution of dimension 1.0 is found using a window having a very large dynamic range of values, on the order of 10^{15} in some cases. In general these windows tend to pile up at the end of the interval, and yield Gabor coefficient distributions in which the

large coefficient is offset to the last time point but not offset in frequency. The instability of these is seen when the attempt to reconstruct the signal in GPS after clipping off the low coefficients fails violently, apparently due to numerical errors. The NAG routines usually reach a happy termination in these cases, confirming that the cost function being used is insensitive to the behavior.

In other cases, the optimizations terminate at low (well under 2.0) dimensions and find windows that resemble the signal except for some usually spiky irregularities. These runs seem to be finding local, not global, minima, and are cause for some concern, because the excessive behaviors of some of the other undesired termination types are not present here. On the other hand, the results are not unusable. Convergence to slightly suboptimum answers may in many cases be almost as good as finding the best answer. Although we would be well advised to understand the phenomena at work here, it is not clear to what extent these results represent problems to solve in the future.

The issues raised by the single signal cases remain when multiple signals are present. There are cases in which the window 'blows up' as described earlier, but still yields a low dimension representation for each signal.

In other respects the results are encouraging. In this light we cite the experiments in which the optimum window was sought for a signal set containing time translates of a decaying exponential. The dimension of such sets was computed for families of exponential windows, and the intuition about those results is that the exponential window is not quite optimum. It was not until the multi-signal optimization code was working that we could look for the optimum in such a case. When we did, we found a window with strong exponential characteristics that stably represented all eight signals. The dimension of the set was similar that found with the best strictly exponential window.

New issues encountered during the multi-signal runs relate to computation time. The effort to compute dimension or its derivatives scales linearly with the size of the signal set, and there is no apparent way to decrease this. Overall run time is somewhat unpredictable despite our understanding of the computational burden of dimension-related evaluations, because the number of evaluation calls is both data- and method-dependent. These observations tend to confirm the suspicion that optimum window analysis is an off-line activity. This is not at all a poor finding, given that many of the envisioned applications of the techniques—e.g., signal detection and classification—might fall into such categories.

Results of the Matlab minimization runs indicate that the choice of algorithm is somewhat data dependent. From the six algorithms tested, two minimized the dimension in all the experiments, and the other four algorithms failed for certain experiments. Therefore the first two algorithms can be considered more reliable, but not necessarily the optimum in all cases.

It was also observed that the number of iterations was reduced by using the gradient if it was available. When the gradient was supplied, the Cubic Interpolation line search method was seen to perform better than Mixed Polynomial Interpolation; fewer function evaluations were performed, but more gradients were evaluated.

The Simplex Search minimization method performed the worst, resulting in a successful minimization in only two out of six experiments. In the two where it did finish, the number of iterations was the largest of all the methods.

The NAG optimization routines have been shown to be capable of determining Gabor representations with minimum dimension. This capability has been demonstrated both for simple functions and more complex real world signals, and for single and multiple signals.

The initial estimate of the biorthogonal is only marginally important in controlling if, and how efficiently, the optimizing routine finds a solution.

In general, the quasi-Newton optimization method employing a user supplied first derivative of the cost function (dimension) yielded the best results. It required fewer iterations to converge than the quasi-Newton method with the finite difference approximation to the first derivative, and ran considerably faster than the Newton method which employed a user supplied second derivative. Each method seemed equally capable of finding solutions.

Linear constraints on the cost function during optimization were helpful when a priori information about the biorthogonal was available. For the more complex transient signals, where the optimum biorthogonal is not known, linear constraints were not helpful.

The best number of frequency and time points (M and N) for a given number of samples depends somewhat on the signal. For signals with detailed features near the beginning, better solutions were generally obtained when $M > N$. For signals without meaningful content near the beginning, $M < N$ generally produced better results.

Solutions were found for the multiple signal experiments where the average dimension was less than if one of the signals were employed as the window function.

Use of the free parameter in the dimension formula (exponent r) was not fully exploited in the tests, but we were able to determine some consequences of varying it. Larger values of r were found to speed convergence in cases that were already convergent at smaller r . Large r tends to emphasize the greatest Gabor coefficient and push the distribution towards the desired shape; consequently, the optimizer sometimes behaves as though its job has become easier, and presents results that appear to have stopped short of reaching the 'eyeball optimum.' The overall conclusion is that low r is the most sensitive case, and may be best for fine tuning a result, while a larger value of r might be used to get to the vicinity of an acceptable solution. Algorithms that automatically adjust r en route according to criteria related to these observations can be imagined for the future.

4.3.1.3 Conclusions

We have shown that it is possible to begin with a collection of signals and find a Gabor representation of these that is maximally concentrated in the sense of the dimension function defined in Phase I. Some of the computed examples show significant differences between a window function found using optimization techniques and a more naively chosen one, some do not. As expected, the techniques developed here are not automatically candidates for every application. Instead, they provide a body of technique that may enhance the ability to carry out a few procedures. Further evaluation must await their application to some tasks.

4.3.2 Extensions Under Current Contract

4.3.2.1 Optimization Tools

The extensions under this contract were a continuation of the nonlinear optimization experiments using a more complex and powerful NAG routine, E04UCF than the previously used algorithm E04KAF. The important properties of the E04KAF routine are displayed below:

- Quasi-Newton algorithm
- Uses analytic gradient (first derivative)
- Builds up surface curvature information (Hessian, or second derivative)
- Incorporates bounds on independent variables

The NAG routine, E04UCF, is more powerful because it allows the user to control more optimization parameters. It performs nonlinear optimization, including bounds on the variables, linear constraints, bounds on the linear constraints, nonlinear constraints, and bounds on the nonlinear constraints. It uses a sequential quadratic programming algorithm in which the search

direction is the solution of a quadratic programming problem (NAG). Its properties are listed below:

- Sequential quadratic programming algorithm
- Minimizes smooth nonlinear function subject to constraints
- Takes analytically specified first partials
- Approximates unspecified first partials by finite differences
- Incorporates linear and nonlinear constraints
- Incorporates *bounds* on independent variables, linear constraints, nonlinear

c

4.3.2.2 Experiments

Numerous experiments were performed with the E04UCF optimization routine to determine average windows for several test cases.

The NAG routine E04UCF was used with several combinations of bounds, linear and nonlinear constraints. Experiments were performed with two signal sets: a single signal, and two signals. The one signal case used a real signal chosen from the first data set of the fault identification section, with a Gabor lattice of $M = 4$, $N = 16$. The two signal case was composed of a rectangular pulse and a decaying exponential, with $M = 4$, $N = 4$.

The bounds and constraints were chosen from the available set of:

- Bounds on the variables
- Linear constraints
- Bounds on linear constraints
- Nonlinear constraints
- Bounds on nonlinear constraints

The choice of these bounds and constraints permitted the control of the minimization process to avoid convergence to unstable results.

Three types of constraints were used and they are described below:

- Type 1: Nonlinear. Restricts exponential growth of biorthogonal signal envelope

Desired constraint is of the form shown in Fig. 4.3-1, but hard to use because of lack of smoothness

$$\max\{|b_1|, |b_2|, |b_3|, |b_4|\} \geq \alpha \max\{|b_5|, |b_6|, |b_7|, |b_8|\}$$

$$\max\{|b_5|, |b_6|, |b_7|, |b_8|\} \geq \alpha \max\{|b_9|, |b_{10}|, |b_{11}|, |b_{12}|\}$$

$$\max\{|b_9|, |b_{10}|, |b_{11}|, |b_{12}|\} \geq \alpha \max\{|b_{13}|, |b_{14}|, |b_{15}|, |b_{16}|\}$$

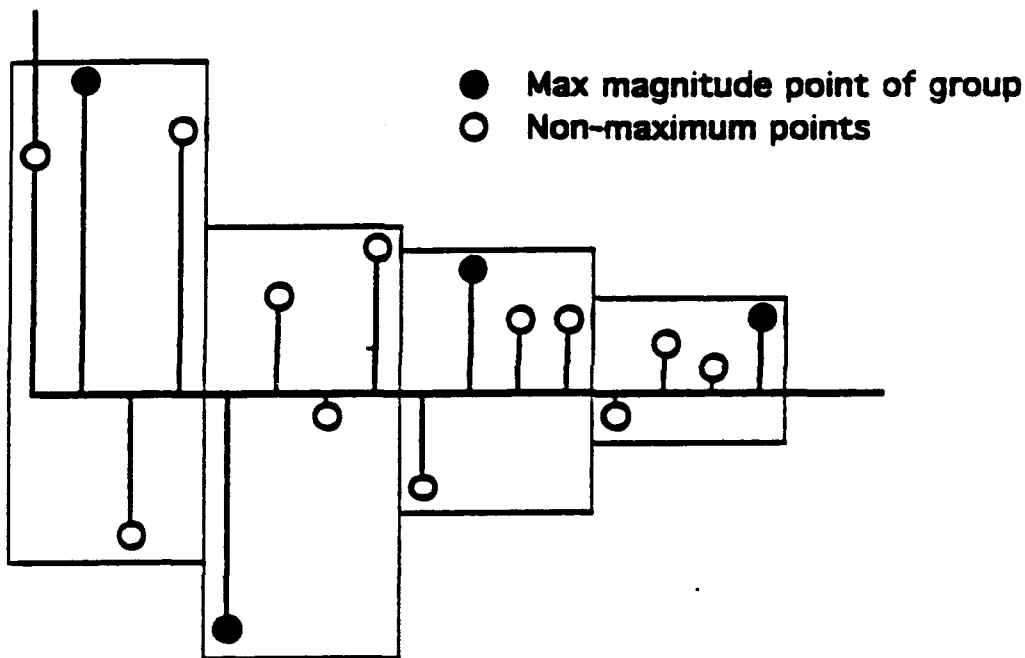


Fig. 4.3-1 Nonlinear constraint that controls exponential growth of the biorthogonal on the basis of the largest magnitude value within each Gabor time slice.

APPLICATION OF GABOR REPRESENTATION TO MILITARY
PROBLEMS(U) ATLANTIC AEROSPACE ELECTRONICS CORP
GREENBELT MD R 5 ORR 2 JUL 93 AFOSR-TR-94-0086

F49620-90-C-0016

NL



END
FILMED
DTIC




AIMM

Association for Information and Image Management



1100 Wayne Avenue, Suite 1100
Silver Spring, Maryland 20910
301/587-8202




Centimeter



Inches



MANUFACTURED TO AIMM STANDARDS
BY APPLIED IMAGE, INC.



This type of constraint was replaced by pointwise comparisons as shown in Fig. 4.3-2:

$$\begin{aligned}
 |b_1| &> \alpha |b_5|, & |b_2| &> \alpha |b_6|, & |b_3| &> \alpha |b_7|, & |b_4| &> \alpha |b_8| \\
 |b_5| &> \alpha |b_9|, & |b_6| &> \alpha |b_{10}|, & |b_7| &> \alpha |b_{11}|, & |b_8| &> \alpha |b_{12}| \\
 |b_9| &> \alpha |b_{13}|, & |b_{10}| &> \alpha |b_{14}|, & |b_{11}| &> \alpha |b_{15}|, & |b_{12}| &> \alpha |b_{16}|
 \end{aligned}$$




Fig. 4.3-2 Nonlinear constraint that controls exponential growth of the biorthogonal on the basis of the magnitudes point-by-point within each Gabor time slice.

- Type 2: Nonlinear. Controls dynamic range of biorthogonal within first Gabor time interval. The need for this can be seen by examining the inversion that creates the window function from the biorthogonal in the matrix Gabor method. Not usable due to non-smoothness. Not needed thus far.

$$\frac{\max_{1 \leq i \leq 4} \{|b_i|\}}{\min_{1 \leq i \leq 4} \{|b_i|\}} \leq \beta$$




Fig. 4.3-3. Nonlinear constraint that controls range within the first Gabor time interval.

- Type 3: Linear: Restrict positive exponential growth of biorthogonal

$$b_1 > \alpha b_{13}, b_2 > \alpha b_{14}, b_3 > \alpha b_{15}, b_4 > \alpha b_{16}$$




Fig. 4.3-4. Linear constraint that controls exponential growth of the biorthogonal on the basis of the largest magnitude value within the first and last Gabor time slice.

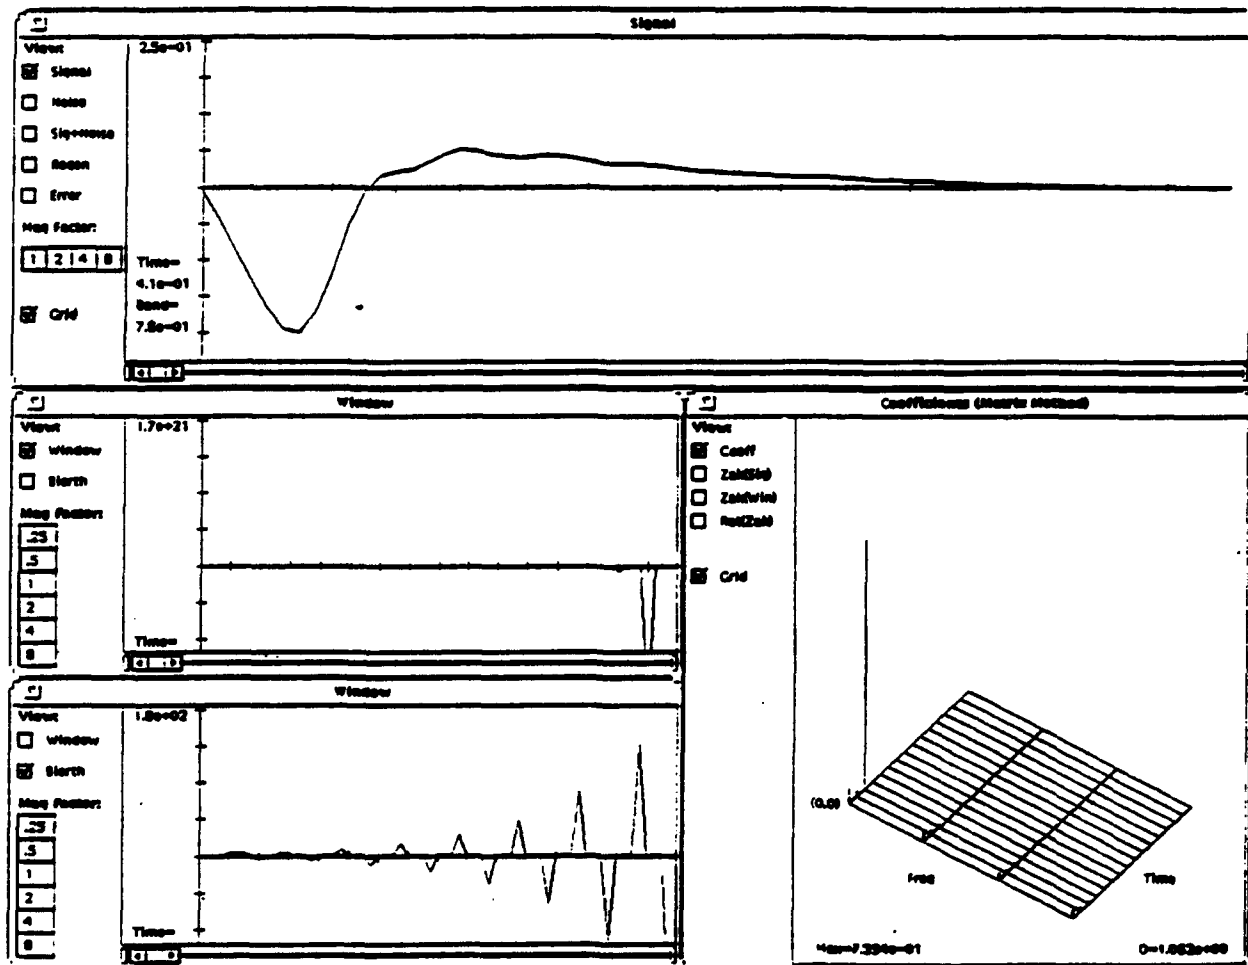
The objective, or cost function, computed the minimum dimension of the signal set for each of the experiments shown in the following figures. For the single signal figures, the initial biorthogonal was a rectangular pulse.

The single signal cases were made with a transient signal from the first data set. Fig. 4.3-5 shows the constraints used, the signal, minimized biorthogonal, the window, the Gabor coefficients, and the dimension was found to be 1.06. This appears to be an excellent minimization result, because the dimension is very close to 1. However, the resulting window is a single negative spike with amplitude of $1.7e+21$, not a useful window. This is the result of convergence to an unstable minimum. This indicates that the minimization process needs to be constrained to avoid convergence to this kind of unstable minimum.

Fig. 4.3-6 shows the same case as in the previous figure, except that the NAG routine E04UCF was used, and nonlinear bounds and constraints were applied. The resultant window was reasonable, and the dimension was still low, at 1.75.

Single signal - Method E04KAF

Bounds on variables	Y
Linear constraints	N
Bounds on linear constraints	N
Nonlinear constraints	N
Bounds on nonlinear constraints	N




$$D = 1.06$$

Fig. 4.3-5. Transient signal - bounds, no constraints

Single signal - Method E04UCF

Bounds on variables	Y
Linear constraints	N
Bounds on linear constraints	N
Nonlinear constraints	Y
Bounds on nonlinear constraints	Y



$$D = 1.75$$

Fig. 4.3-6. Transient signal - bounds, and nonlinear constraints

Several experiments were performed with a two signal set: a rectangular pulse, and a delayed exponential. For all of the two-signal minimization experiments, a random initial biorthogonal was chosen. This initial biorthogonal and its corresponding window are shown in Fig. 4.3-7.

Before attempting to find an optimum window, each of the two signals were used as a window in the computation of Gabor coefficients, as shown in Fig. 4.3-8 and 4.3-9. In Fig. 4.3-8, the pulse is used as a window, which happens to be the same as the biorthogonal. The lower coefficient set reproduces the pulse, and the upper set reproduces the exponential. The average of the two dimensions is 1.98. Similarly, in Fig. 4.3-9, the exponential used as a window results in a average dimension of 1.503. What we are looking for is an optimum window that produces a minimum average dimension. Figure 4.3-10 shows the optimization result using the NAG routine E04KAF, which uses bounds on the variables only. The resulting average dimension is 1.496, very close to the previous case.

The next experiment (Fig. 4.3-11) used the E04UCF method with bounds on the variables and linear constraints with bounds. The four linear constraints used were the Type 3 constraints mentioned earlier. The result was $D = 2.00$, showing that this type of constraint did not succeed in decreasing the average dimension.

Figure 4.3-12 shows the result of using bounds on the variables and twelve Type 1 nonlinear constraints with bounds. The average dimension is lower, at 1.887.

The value of α , the coefficient in the Type 1 constraint equation was then changed from 0.5 to 1.01. The previous case was run again, resulting in Fig. 4.3-13. The minimized biorthogonal was constrained as expected, but the average dimension was higher. The same case was run again with $\alpha = 1.02$, with the result shown in Fig. 4.3-14. The average dimension is now lower, but the biorthogonal has changed to nearly a rectangular pulse. Apparently, the optimization algorithm is very sensitive to changes in the constraint function.

- Initial biorthogonal chosen randomly (upper panel)
- Corresponding window (lower panel)


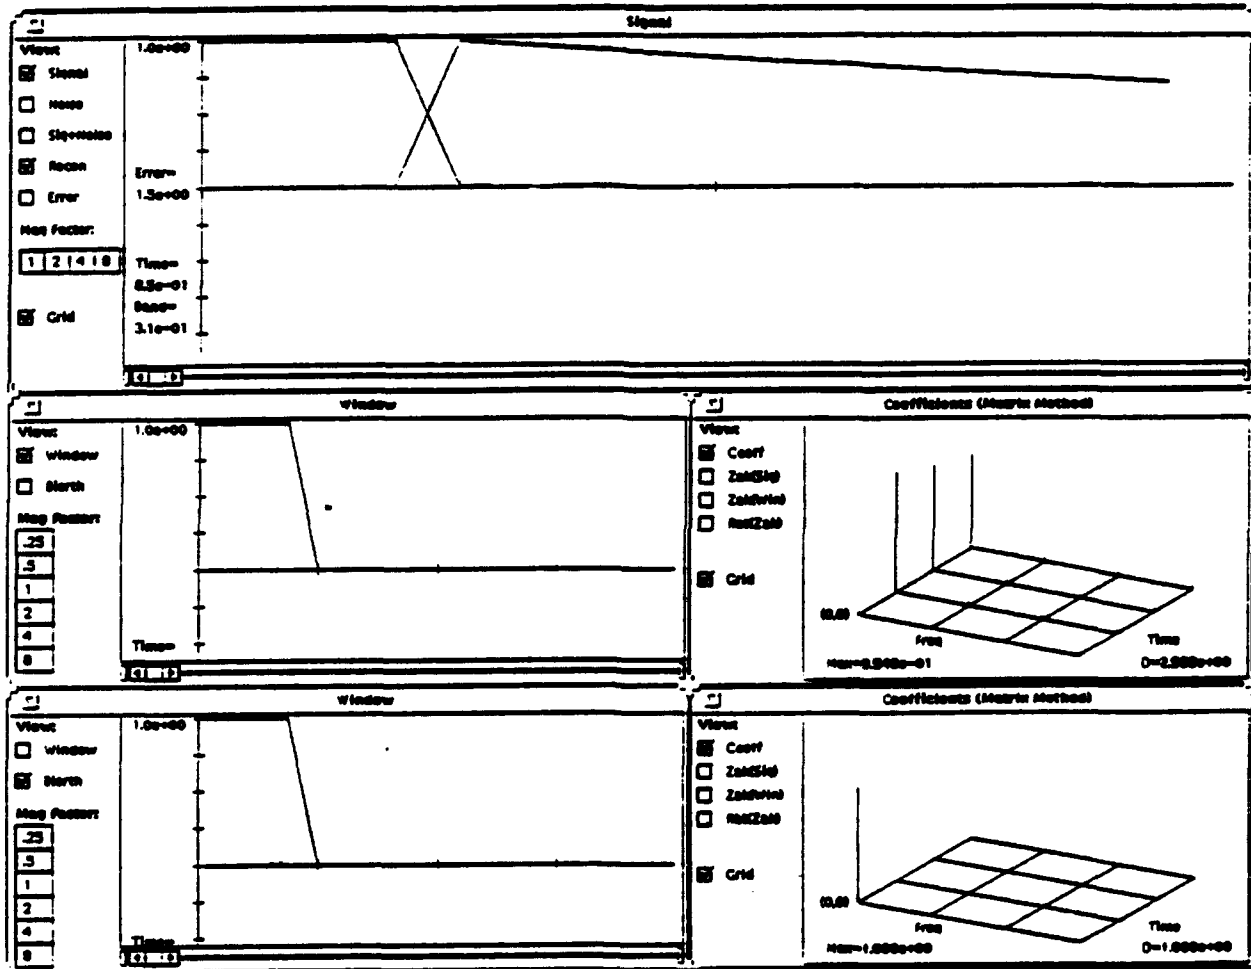


Fig. 4.3-7. Initial biorthogonal and corresponding window.

Two signals - No optimization

Window equals rectangular pulse

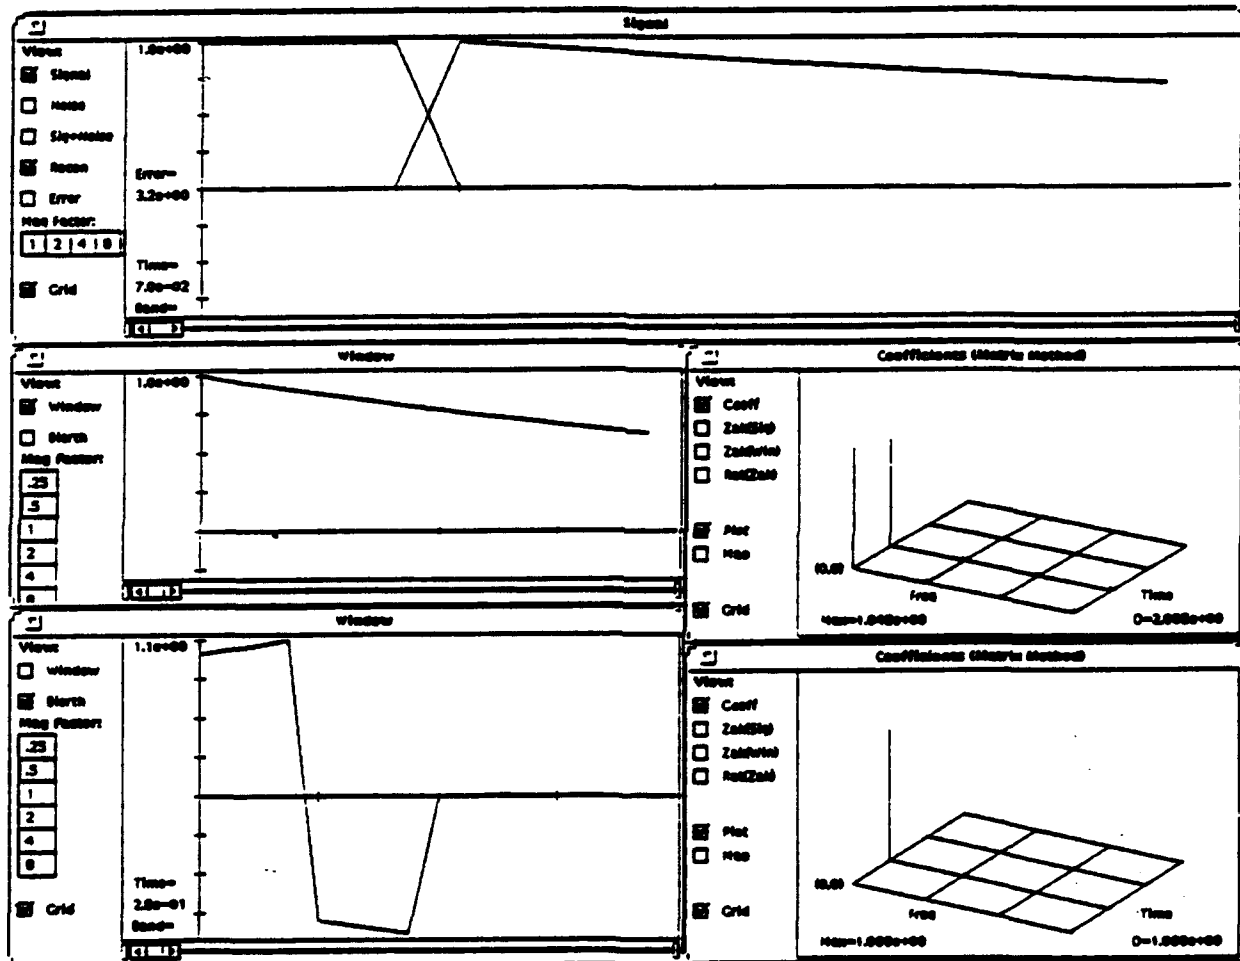


$$D_{av} = 1.98$$

Fig. 4.3-8. Two signals - no optimization, window equals rectangular pulse.

Two signals - No optimization

Window equals exponential pulse

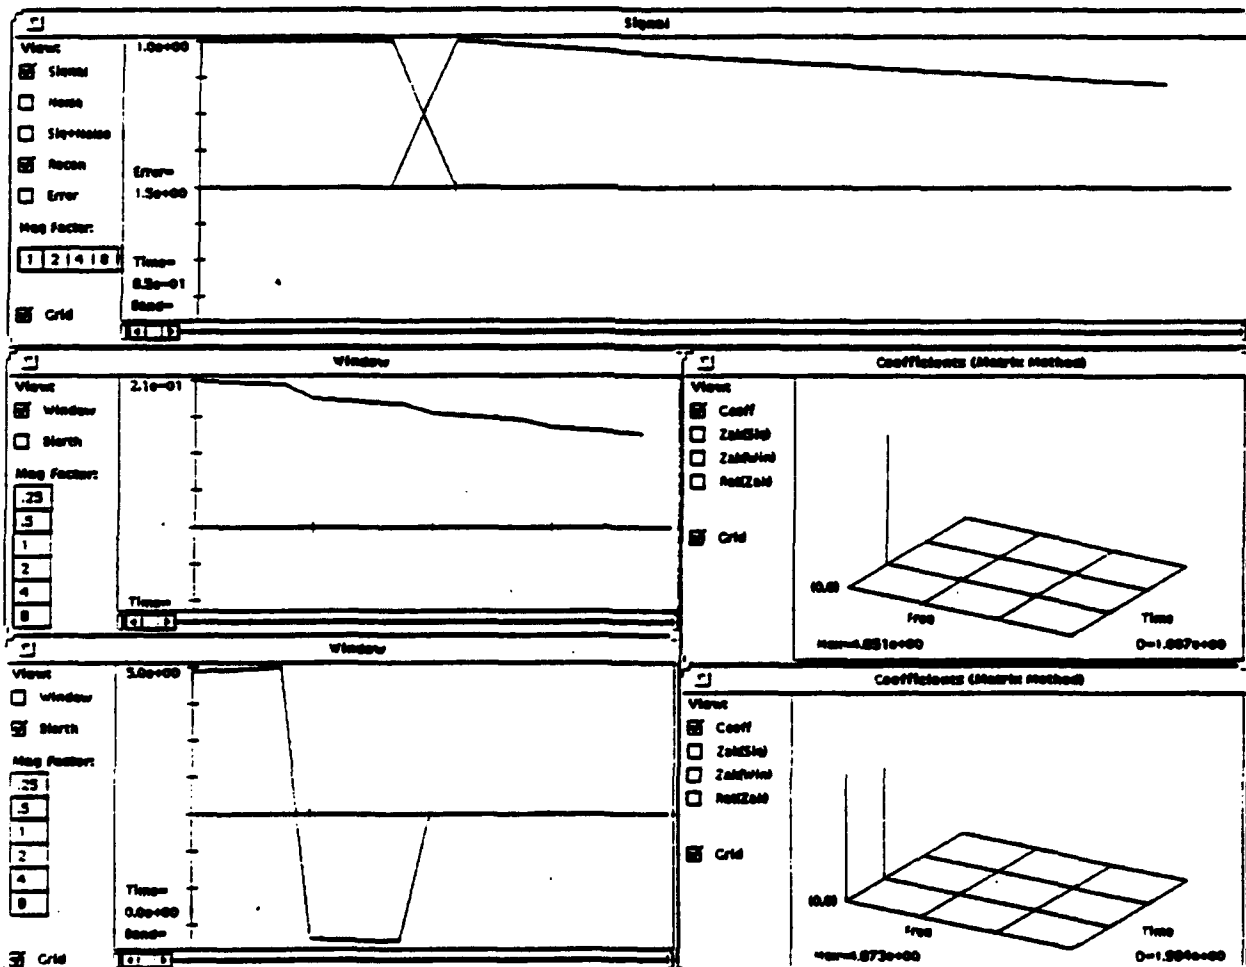


$$D_{av} = 1.503$$

Fig. 4.3-9. Two signals - no optimization, window equals exponential pulse.

Two signals - Method E04KAF

Bounds on variables	Y
Linear constraints	N
Bounds on linear constraints	N
Nonlinear constraints	N
Bounds on nonlinear constraints	N

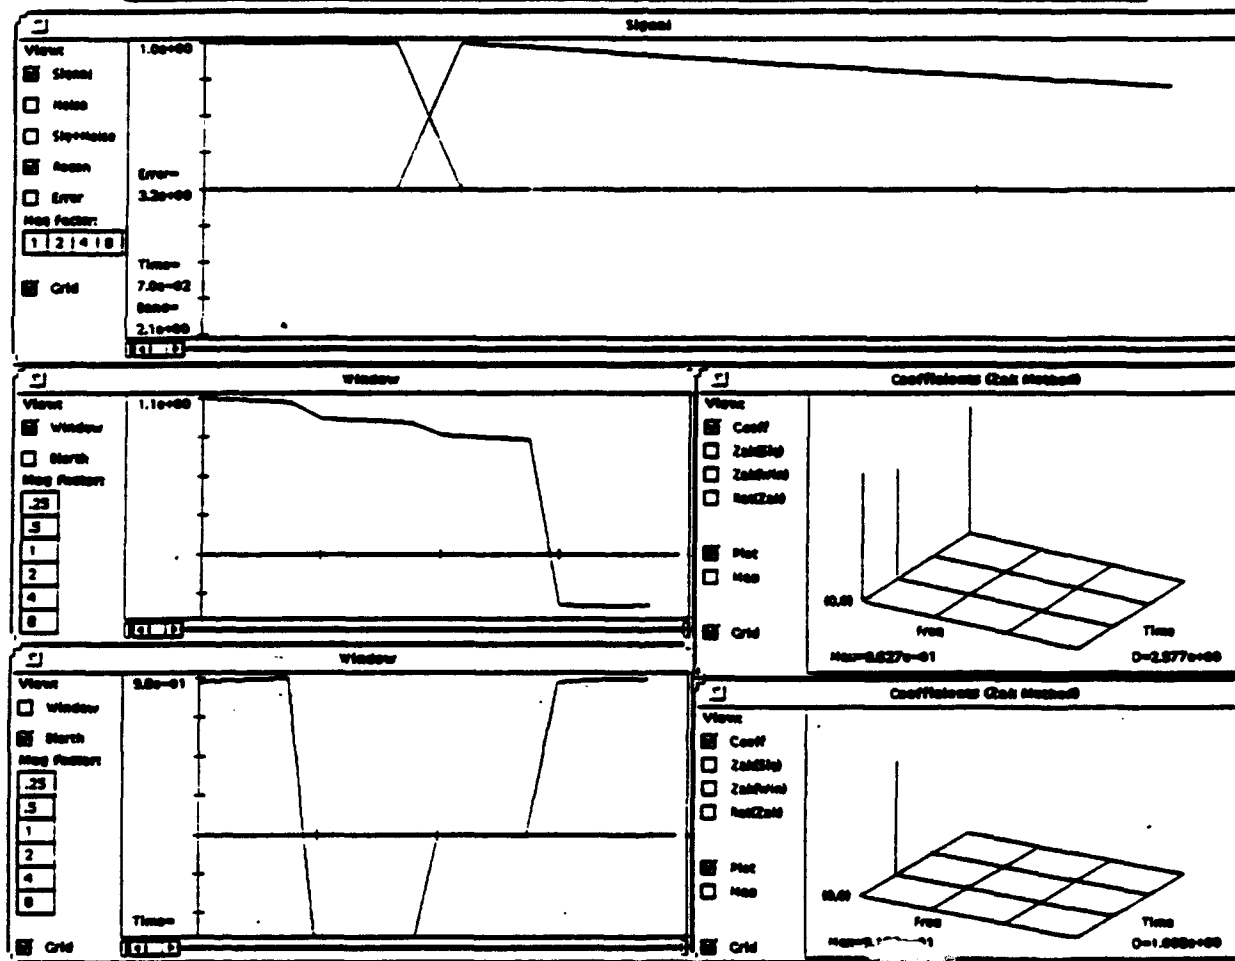


$$D_{av} = 1.496$$

Fig. 4.3-10. Two signals- bounds, no constraints.

Two signals - Method E04UCF

Bounds on variables	Y
Linear constraints	Y
Bounds on linear constraints	Y
Nonlinear constraints	N
Bounds on nonlinear constraints	N



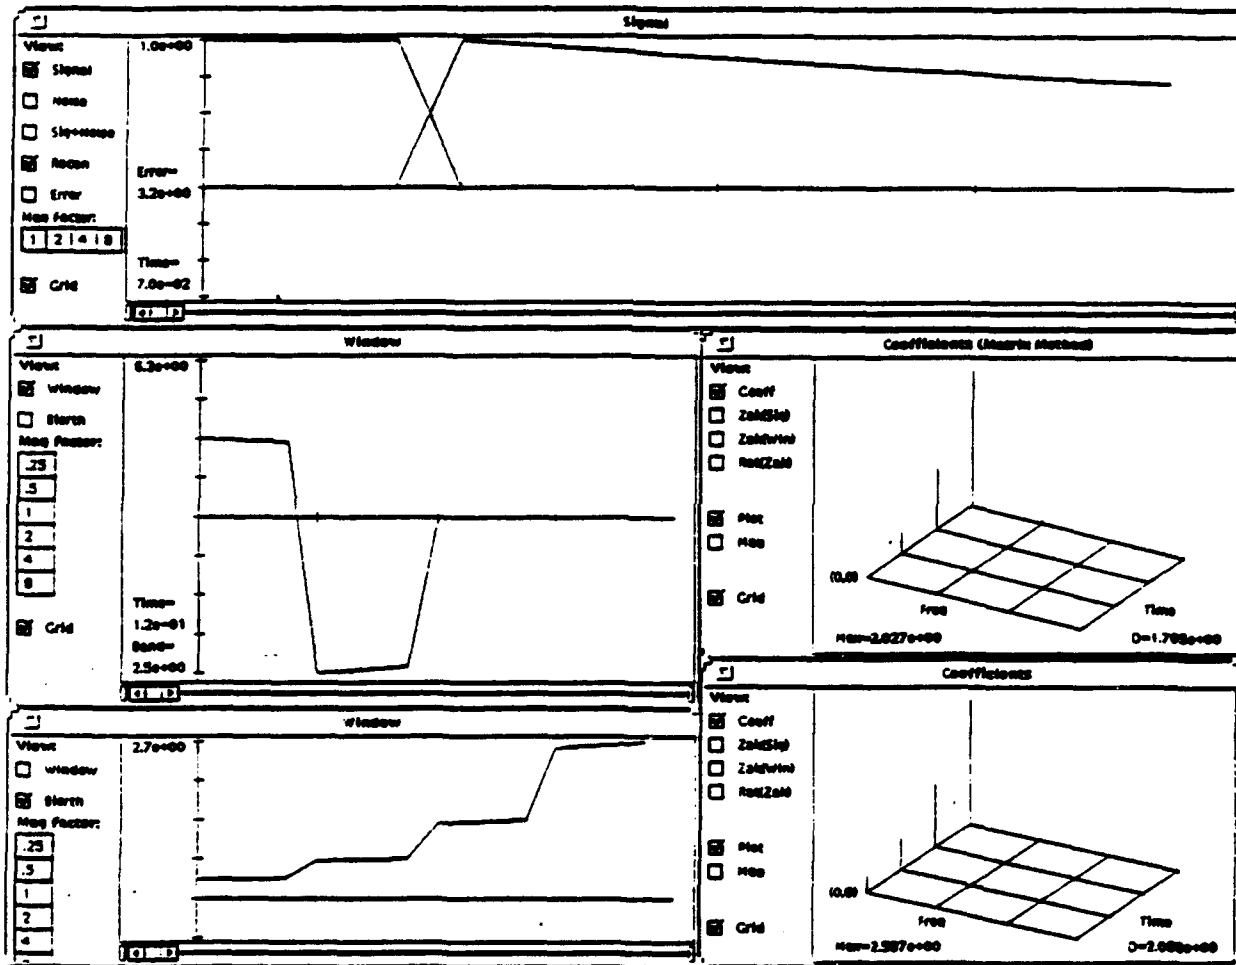
$$D_{av} = 2.00$$

Fig. 4.3-11. Two signal case - bounds, linear constraints.

Two signals - Method E04UCF ($\alpha = .5$)

- 12 constraints

Bounds on variables	Y
Linear constraints	N
Bounds on linear constraints	N
Nonlinear constraints	Y
Bounds on nonlinear constraints	Y

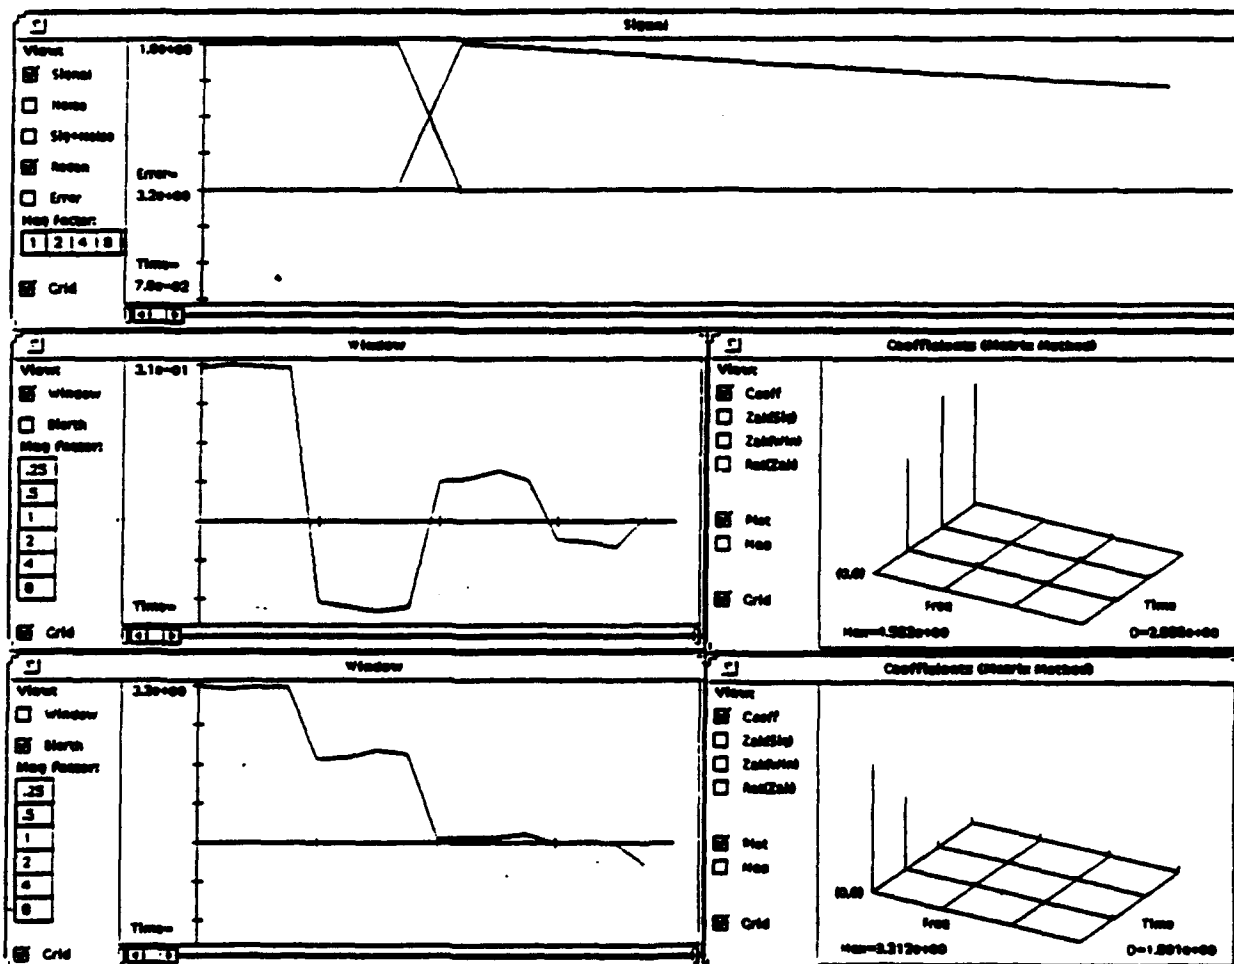


$$D_{av} = 1.887$$

Fig. 4.3-12. Two signal case - bounds, nonlinear constraints.

Two signals - Method E04UCF ($\alpha = 1.01$)

Bounds on variables	Y
Linear constraints	N
Bounds on linear constraints	N
Nonlinear constraints	Y
Bounds on nonlinear constraints	Y

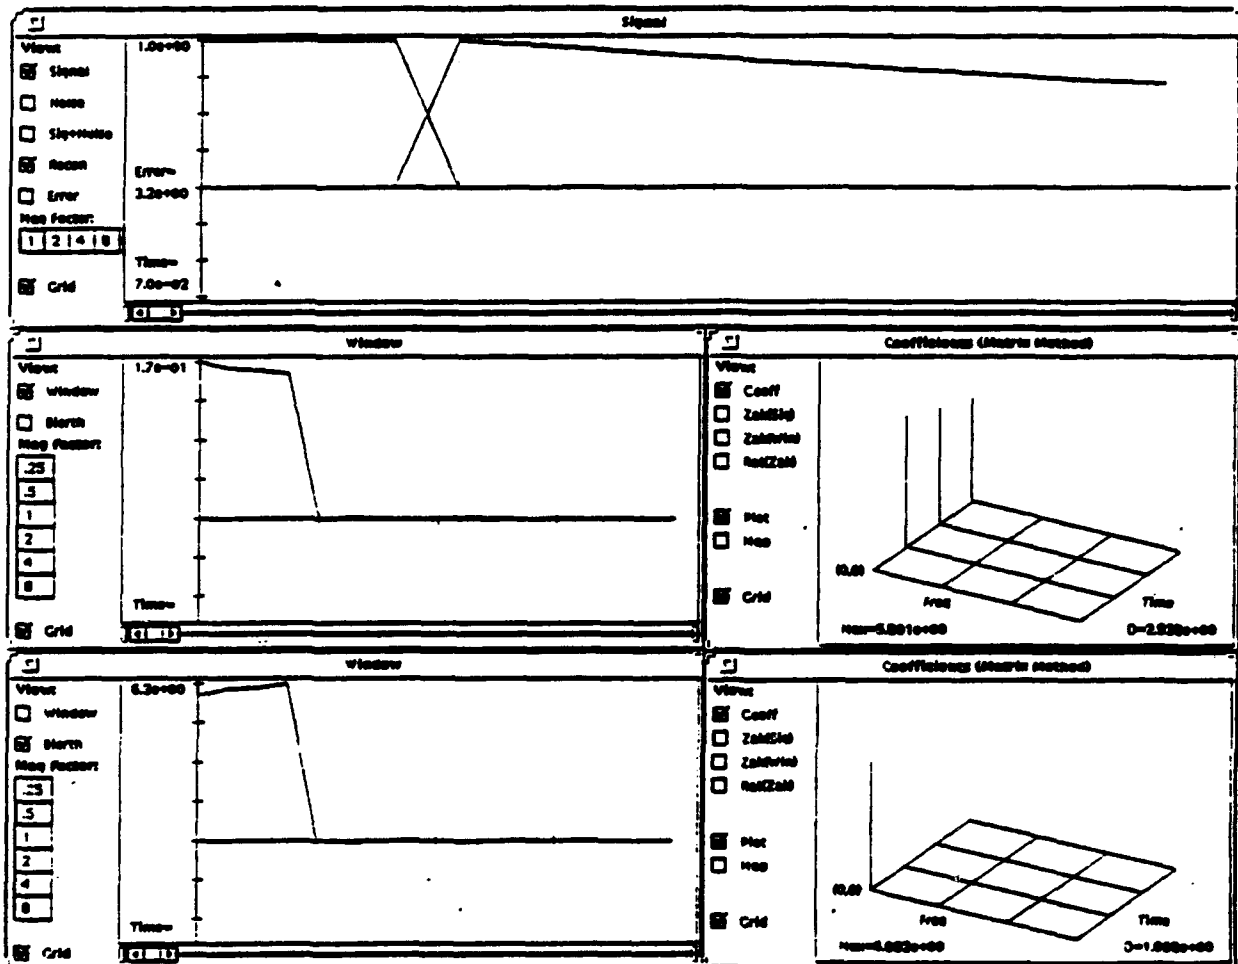


$$D_{av} = 2.35$$

Fig. 4.3.13. Two signal case - bounds, nonlinear constraints.

Two signals - Method E04UCF ($\alpha = 1.02$)

Bounds on variables	Y
Linear constraints	N
Bounds on linear constraints	N
Nonlinear constraints	Y
Bounds on nonlinear constraints	Y



$$D_{av} = 1.97$$

Fig. 4.3-14. Two signal case - bounds, nonlinear constraints.

4.3.2.3 Observations

With these additional optimization runs, we have managed to minimize the dimension where the biorthogonal function previously did not converge to a reasonable state. The addition of nonlinear constraints and bounds, along with the sequential quadratic programming minimization algorithm, has given us much more control over the optimization process. It was observed that the minimization can be very sensitive to changes in the constraints. This was seen in the significant change in the resulting window function resulting from a small change in the value of alpha in the nonlinear constraint equation.

4.3.3 Future Work

The work presented above makes a positive extension of the results originally obtained under the SBIR and provides further justification of the ideas that first lead to this line of endeavor. The task remains unfinished, however. Research to date has not completed the task of determining the best constraints to use in optimizing windows, and the bugs in using the optimization software have not yet been completely worked out. Work has been slowed somewhat by this situation. Using simulated annealing for optimization has also been suggested, and AAEC is currently looking into available software packages for that purpose.

In summary, the effort so far has proved in principle most of the supporting concepts, but has been insufficient to transition the work into the applications arena as yet. AAEC sees particular promise for this technology in certain applications areas, and is planning to propose effort in those areas. A key area is automatic target recognition (ATR). Initial work in applying the Gabor transform to problems within that discipline is reported elsewhere in this document. Machine-aided recognition problems have the feature that searching for objects can be enhanced in circumstances where shape characteristics of the objects are partially known in advance, either through a priori knowledge or data-aided algorithms. For example, in signal analysis, the Gabor transform is particularly adept at finding features having a common envelope.

To maximally exploit such a circumstance, extraction of the analysis window from the data looms important. Given a large body of data such as that often encountered in an ATR problem, use of the data to drive the analysis functions seems wise as a measure to cut the amount of blind search, especially in view of findings that allegedly more "robust" tools such as the Wigner distribution can create artifacts through nonlinear processing if not used carefully. The role for optimum Gabor windowing in this scheme is clear, and as a result it appears that the best way in

which to continue the line of work discussed above is to do it within the context of an application area such as ATR.

The research is at the point where it could profit from the interaction with real data as an aid in algorithm development/refinement. AAEC anticipates proposing a body of work of this nature as a logical follow-on the work performed in both this contract and the cited SBIR's.

Atlantic Aerospace Electronics Corporation

REFERENCES

- AAEC91(1) *Application of Gabor Representation to Military Problems*, Atlantic Aerospace Electronics Corporation Technical Report, Greenbelt, MD, 4 February 1991.
- AAEC91(2) *Determination of Minimum Dimension Gabor Representations for Signal Detection and Characterization*, Atlantic Aerospace Electronics Corporation Technical Report, Greenbelt, MD, 30 Jan., 1991.
- AAEC92(1) *Application of Gabor Representation to Military Problems*, Atlantic Aerospace Electronics Corporation Technical Report, Greenbelt, MD, 4 February 1992.
- AAEC92(2) *Development Of A Minimum Dimension Gabor Representation Facility*, Atlantic Aerospace Electronics Corporation Technical Report, Greenbelt, MD, 30 September 1992.
- BALA1 Balart, R. and R. S. Orr, "Computational Accuracy and Stability Issues for the Finite, Discrete Gabor Transform," *Proc. IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, 1992.
- BALA2 Balart, R., "Matrix Reformulation of the Gabor Transform", *Opt. Eng.* 31 (6), pp. 1235-1242, 1992.
- KOHO T. Kohonen, J. Kangas, J. Laaksonen, *SOM-PAK - The Self-Organizing Map Program Package (Version 1.2)*, November, 1992.
- NAG National Algorithms Group Users Manual, Mark 15, Vol. E.
- SWEE J. Sweeney, R. Orr, R. Balart, D. Buchanan and A. Humen, "Experiments in Dimensionally-Optimum Gabor Representations," *Proc. IEEE ICASSP93*, Vol. III, pp. 265-268.
- ORR Orr, R. S., "Dimensionality of signal sets," *Proc. SPIE*, Vol. 1565, pp. 425-446, 1991.
- SERR J. Serra, *Image Analysis And Mathematical Morphology*, Vol. 1, Academic Press Ltd., 1982.
- SHAN Shannon, C. E., "Communication in the Presence of Noise," *Proc. IRE*, Vol. 37, No. 1, pp. 10-21, January 1949.
- TOLI Tolimieri, R. and A. Brodzik, *Convergence of Discrete Gabor Expansions for General Window Functions*, preprint, 1992.
- WEXL Wexler, J. and S. Raz, "Discrete Gabor Expansions," *Signal Processing*, Vol. 21, pp. 207-220, 1990.

APPENDIX A

STATEMENT OF WORK

AAEC will perform the following work items, segregated by year.

Year 1:

- **Proposed Theoretical Developments:** Study methods to improve the stability and accuracy of the Gabor coefficient computations; perform error analyses to understand the accuracy of approximation achieved by truncated, finite Gabor expansions.
- **Software Development:** Develop software code on a GFE Aspen computer to execute the new algorithms for Gabor analysis. This includes full debugging and achieving operational status.
- **Comparative Analysis of Competing Methods:** Investigate the utility of competing methods, including other double series representations, Wigner distributions, non-abelian harmonic analysis, etc., for the problems under attack.
- **Potential Applications:** Begin to investigate potential application of Gabor representations to military problems, and use the tentative results to guide the needed theoretical developments. Test problems to be chosen from those in the proposal.

Year 2:

- **Potential Applications:** Continue to investigate potential application of Gabor representations to military problems, and use the tentative results to guide the needed theoretical developments. Test problems to be chosen from new areas discovered during first-year investigations and discussions with interested parties at DARPA and elsewhere.
- **Software Upgrades:** Maintain existing software and add new capabilities as needed to improve ability to model potential applications.

Year 3:

- **Continuation of Year 2 work items.**

Full Statement of Work disclosed to DARPA and available upon request.

APPENDIX B

ACCURACY AND STABILITY PAPER

Atlantic Aerospace Electronics Corporation

COMPUTATIONAL ACCURACY AND STABILITY ISSUES FOR THE FINITE, DISCRETE GABOR TRANSFORM

Rogelio Balart and Richard S. Orr

Atlantic Aerospace Electronics Corporation, 6404 Ivy Lane, Suite 300, Greenbelt, MD 20770 USA

ABSTRACT

A Gabor expansion may employ highly nonorthogonal basis functions and consequently inherit accuracy and stability problems due to near-singularity when computed digitally. Two methods to discretize the Gabor transform are studied from the viewpoint of controlling numerical properties: (i) a Zak transform-based method and (ii) a matrix method. Theoretical issues relating to the singular behavior of each are cited, and stabilization techniques are proposed. We demonstrate the validity of each technique with results of numerical experiments, concluding that stability and accuracy can usually be achieved in a digitally implemented Gabor transform by proper choice of algorithm and stabilization mechanism.

1. INTRODUCTION

A Gabor expansion of an $L^2(\mathbb{R})$ function is a doubly infinite series of the form

$$f(t) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{m,n} w_{m,n}(t) \quad (1)$$

in which the $\{a_{m,n}\}$ are the *Gabor coefficients*, and the $\{w_{m,n}\}$ are time and frequency translations of a *window function*, $w(t)$, that are the basis of the expansion. The properties of the expansion are highly dependent upon the behavior of the window with respect to the time-frequency lattice $\{(f_m, i_n)\} = \{(m/T, nT)\}$ that supports the basis. The parameter $T \geq 0$ mitigates the time and frequency resolution of the transform. Many of these properties are readily expressed in terms of the Zak transform of w .

$$Zw(v, \tau) = \sum_{k=-\infty}^{\infty} w(kT + \tau) \exp(-j2\pi k v T). \quad (2)$$

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Air Force Office of Scientific Research under contract No. F49620-90-C-0016. The United States Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

If the modulus of the Zak transform of w , $|Zw|$, is constant almost everywhere, the $\{w_{m,n}\}$ are orthogonal and well behaved. When $|Zw|$ has a small dynamic range (small ratio of its maximum and minimum values), the expansion is nonorthogonal but still quite stable. The difficult case is when Zw has a zero within $0 \leq t < T$, $0 \leq n < 1/T$. The formal expression for the Gabor coefficients,

$$a_{m,n} = \int_0^T \int_0^{1/T} \frac{Zf(v, \tau)}{Zw(v, \tau)} \exp(j2\pi(nvT - m\tau/T)) dv d\tau, \quad (3)$$

does not then necessarily represent the Fourier coefficients of an $L^2(\mathbb{R}^2)$ function, and it implies an expansion for which the above equation does not always yield a reliable basis for numerical evaluation of the $\{a_{m,n}\}$.

For computation by the Zak method, (3) must be replaced by a finite, discrete equivalent. Any algorithm for this has the result

$$a_{m,n} = \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} \frac{Zf(v_q, \tau_p)}{Zw(v_q, \tau_p)} \exp[j2\pi(nv_qT - m\tau_p/T)], \quad (4)$$

Sampling and periodization can be applied in a way that retains much of the correspondence to the continuous time expansion [1, 2]. In this paper we consider transforms made discrete by sampling and truncation.

Most of the continuous time accuracy/stability issues carry over to the discrete case. In particular, it is easy to see in (3) the impact of a zero of Zw . Even if Zf and Zw have a coinciding zero, the ratio Zf/Zw may not be well defined. When $|Zw|$ makes a close approach to zero at some grid point, the ratio can become very large there, and dominate the expression for the coefficients. Ways to control transforms with Zak zeros are given in section 2.

A second approach is to treat the sampled equations for the Gabor transform as a matrix and solve for the coefficients by matrix inversion. The structure of the matrix is such that the inversion can be performed efficiently [3, 4]. For reasons of completeness, we will now introduce the rudiments of the matrix formulation. If one starts with the Gabor transform given by (1),

uniformly samples f at every $\Delta t = T/M$ ($t_k = k(T/M)$), $k=1, 2, \dots, MN=P$, and truncates the expansion after P points, one obtains the representation

$$f(t_k) = \sum_{n=0}^{M-1} \sum_{m=0}^{N-1} a_{nm} w(t_k - nT) \exp(j2\pi m t_k / T) \quad (5)$$

which can be written in matrix form as

$$F = (W E) A, \quad (6)$$

where F and A are $P \times 1$ column vectors corresponding to the signal and the (reordered) coefficients respectively, and W and E are $P \times P$ matrices corresponding to the values of the window and the exponentials associated with the rotations around the unit circle. As it turns out, the matrix W is block lower triangular and each of the sub-blocks is diagonal. The sub-blocks which lie on the diagonal are all equal and correspond to the values of the window over the first M points, and similarly all the sub-blocks on the m th off-diagonal are equal and contain the M values of the window after skipping over $LM - 1$ points.

The matrix E also has a particularly nice form. It is block diagonal, all of the sub-blocks are equal, and the entries of the sub-blocks are the entries of the Fourier rotation matrix. Consequently, E is invertible analytically. With this in mind, we can solve (6) for the unknown vector of Gabor coefficients and we obtain

$$A = E^{-1} W^{-1} F. \quad (7)$$

Since the inverse of E is analytically calculated and well behaved, the stability issues are solely associated with the invertibility of W . The structure of the inverse of W and the resulting stability issues will be discussed in section 3.

2. ZAK METHOD EXPERIMENTS

2.1 Theory

The stability of a continuous time Gabor representation can be gauged by the size of the ratio

$$C = \frac{\max_{0 \leq v < 1/T, 0 \leq \tau < T} |Zw(v, \tau)|^2}{\min_{0 \leq v < 1/T, 0 \leq \tau < T} |Zw(v, \tau)|^2}, \quad (8)$$

which acts as a condition number of the transform. This number is also the frame bound ratio of the transform [5].

When implemented by a discrete Zak method, the Gabor transform is characterized by the same ratio sampled on the time-frequency grid,

$$v = n/NT, \tau = mT/M; 0 \leq n < N-1, 0 \leq m < M-1. \quad (9)$$

If Zw has a zero on the grid, the condition ratio will be infinite and the transform noninvertible. Restricted to (9), C need not be infinite when (8) is, letting the discrete transform possibly be more stable than the continuous.

Three methods of stabilizing the discrete transform are illustrated. We start with a discrete domain window having a Zak transform zero on the grid. In the first method, we move the zero of the Zak transform off the grid by a subtle change in the shape of the window function. The second method is to apply a time and/or frequency translation to the original window that relocates the Zak zero according to the formula

$$Z\{f(t - \tau_0) e^{j2\pi v_0(t - \tau_0)}\} = e^{j2\pi v_0(\tau - \tau_0)} Zw(v - v_0, \tau - \tau_0). \quad (10)$$

A third method is to retain the window function but change to a new grid on which the zero no longer lies.

2.2 Experiments

The window function employed for the experiments is

$$w(t) = t \exp(-at) u(t), \quad (11)$$

where a is a positive decay constant and u is the unit step function at the origin. The Zak transform of (11) is

$$Zw(v, \tau) = \frac{e^{-a\tau} \left\{ \tau + (T - \tau) e^{-a + j2\pi v T} \right\}}{\left[1 - e^{-a + j2\pi v T} \right]^2}, \quad (12)$$

which has a single first-order zero at the point

$$(v^*, \tau^*) = \left(\frac{1}{2T}, \frac{T}{1 + e^{aT}} \right). \quad (13)$$

Notice that the zero lies on a frequency grid point when the number of frequency points M is even, but the time location of the zero is arbitrary within $(T/2, T)$.

For an N by M Gabor transform, the values of a lying on the grid are obtained from (13):

$$a = \frac{1}{T} \ln \left(\frac{M}{m} - 1 \right); \quad 0 < m < M/2. \quad (14)$$

The following experiments use $N = 16$, $M = 16$. Fig. 1 shows the window, and Fig. 2 shows the Zak transform for the choice $a = -1.94591$. Because of single-precision roundoff in the calculation, the condition number (8) is 8.523×10^{15} and not infinite at this point. If we choose the decay constant according to the formula

$$a = \frac{1}{T} \ln \left(\frac{M}{m + 1/2} - 1 \right), \quad (15)$$

the Zak zero will lie midway between grid points. For the choice $a = -1.68640$, we find the Zak transform shown in Fig. 3 with condition number $C = 4.175 \times 10^2$.




Fig. 1 - Analysis window




Fig. 2 - Zak transform of the window. The transform has an analytical zero on a time frequency grid point. The condition number is displayed on the lower left hand corner.

The behavior of a transform having a Zak zero near the grid is quite sensitive to the choice of a . For example, choosing $a = -1.94590$, a change in the 5th significant digit, changes the condition number by 4 orders of magnitude to 1.618×10^{11} . When the zero is placed midway between grid points the sensitivity is much less, for example, both $a = -1.68640$ and $a = -1.68641$ yield $C = 4.175 \times 10^2$ to four significant figures.

Fig. 4 illustrates one application of our second method: the window having the on-grid zero is time-delayed by $T/2M$ (one-half sample) before sampling, and as a result the condition number (8) of the Zak transform becomes $C = 2.513 \times 10^7$.

Experiments showing the third method of relocating the zero of the Zak transform by means of a grid change are not included due to space limitations




Fig. 3 - Zak transform of the analysis window, where the analytical zero has been moved off the grid by adjusting the time constant. The condition number is displayed on the lower left hand corner.




Fig. 4 - Zak transform of the window after the analytical zero has been moved off the grid by delaying the window arrival time. The condition number is shown on the lower left hand corner.

3. MATRIX METHOD EXPERIMENTS

3.1 Theory

As we mentioned earlier, the stability of a discrete time matrix Gabor representation depends only on the inverse of the window function which we will denote by B due to the biorthogonality relationship between B and W . The form of B is identical to the form of W , and its entries are given by the recursive relation

$$B_l = -W_0^{-1}(W_l W_0^{-1} - W_{l-1} B_1 - \dots - W_1 B_{l-1}) \quad (17)$$

where the subscript $l = 0, 1, \dots, N-1$ denotes the distance of the block from the main diagonal. Clearly the only part of W that gets inverted is the block associated with the first M points of the window function, and therefore, a necessary condition for invertibility is that the window be nonzero everywhere in that region. Even if that condition

is satisfied. It does not imply that the matrix B will be well behaved. Notice that as we proceed further down the off diagonals, i.e. as l gets larger, the last term in (17) will be multiplied by the l^{th} power of the inverse of W . This essentially means that comparatively small values of the window over the first M points can lead to values at the tail end of the biorthogonal which are undesirably large. A measure of the stability of the window can at least empirically be characterized by the biorthogonal having no small values over the first M points, compared to the values over the rest of the window values.

3.2 Experiments

The window function used for the matrix experiments is the same as the window used for the Zak method experiments with the time constant corresponding to placing the zero of the Zak transform on one of the grid points. If we choose this window as is, we know a-priori that the biorthogonal will be infinite, since our window function has an analytical zero at the origin, and this violates our initial necessity condition for invertibility. This problem can, of course, be alleviated. If we shift our window to the left by $2T/M$, i.e. two points, the window no longer has a zero on the interval containing the first M points, and the biorthogonal is well behaved as is shown in Fig. 5. The values over the first M points are not significantly smaller than the rest of the values, and the set of values over the whole interval are bounded by, in this case, 13.0. One of the other ways to obtain a nice biorthogonal is to increase the amplitude of the original window by a constant factor, i.e. add a value of 0.2 to every point of the window. The biorthogonal corresponding to this window has values that are bounded by 23.0, and it is displayed in Fig. 6.




Fig. 5 - Biorthogonal to the analysis window whose zero has been shifted two time samples to the left. The maximum value of the biorthogonal is given on the upper left hand corner of the axis.

4. SUMMARY

The initial motivation that led to the development of the matrix representation of the Gabor transform was to

alleviate some of the problems associated with the Zak transform of the window having a zero somewhere in the region of interest. As we have shown in the above experiments, the existence of a second method to calculate the coefficients does not guarantee a stable coefficient set.




Fig. 6 - Biorthogonal to the analysis window after adding a constant value of 0.2 to every window sample. The maximum value of the biorthogonal is given on the upper left hand corner of the axis.

One has to resort to some further knowledge of the structure of the window function in order to obtain results that are not sensitive to small perturbations. It is the belief of the authors that by appropriately tailoring the window, the instabilities associated with one or both of the methods can be minimized, and at least one of the methods will give satisfactory results.

REFERENCES

1. Orr, R. S., "Derivation of the Finite Discrete Gabor Transform by Periodizing and Sampling," accepted for publication, *Signal Processing*.
2. Orr, R. S., "Finite Discrete Gabor Transform Relations Under Periodizing and Sampling," *Proc. IEEE Time-Frequency Time-Scale Symposium*, 1992.
3. Balart, R., "Matrix Reformulation of the Gabor Transform", *Proc. SPIE Adaptive Signal Processing Symposium* 1565, 1991.
4. Balart, R., "Matrix Reformulation of the Gabor Transform", *Opt. Eng.* 31 (6), pp 1235-1242, 1992.
5. Daubechies, I., "The Wavelet Transform, Time-Frequency Localization and Signal Analysis", *IEEE Trans. IT* 36 (5), pp 961-1005, 1990.

APPENDIX C
PERIODIZED GABOR PAPER

Derivation of the Finite Discrete Gabor Transform by Periodization and Sampling

Richard S. Orr

Atlantic Aerospace Electronics Corporation
6404 Ivy Lane, Suite 300
Greenbelt, Maryland 20770
(301) 982-5215

Submitted to *Signal Processing* 16 August 1991

Revised 21 April 1992

2nd Revision 15 December 1992

Abstract.¹ Recent interest in the Gabor transform for time-frequency signal analysis can be attributed in part to increased knowledge about the accuracy, stability, and complexity of algorithms for computing the transforms. Behavior of the computations depends on, among other things, the manner in which the continuous-parameter equations are made discrete and finite. The most straightforward means, truncating the time functions to compact support and sampling, relinquishes some control and blurs the relationship of the discrete equations to the original transforms. A more satisfying discretization and finitization process that preserves relations to the continuous parameter case is found by periodization and sampling, the same method used to obtain the finite, discrete Fourier transform from the Fourier integral. By this method we derive the finite, discrete Gabor transform equations from their continuous parameter counterparts, in the process explicitly exhibiting the aliasings that permit one periodic sequence to be the finite, discrete Gabor transform of the other. By examining the various forms in which the Gabor equations can be expressed, we discover how the input, window, biorthogonal function, Gabor coefficients and Zak transforms map under periodization and sampling.

¹This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the Air Force Office of Scientific Research under contract No. F49620-90-C-0016. The United States Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

1. Introduction

The Gabor transform has recently attracted a certain interest in time-frequency signal processing because of its ability to portray time history of the frequency content of a signal or image by representing the data as a superposition of uniform time and frequency translates of a window function. Through choice of the window, a Gabor representation can exploit a priori information about the signal that is useful in problems of signal detection, feature extraction, and identification.

To some extent, current interest in Gabor technique is fueled by successful developments related to computing the transforms. Since Bastiaans first considered the question of solving for the Gabor coefficients, given the input function and the window [9], it has been recognized that this computation can be ill-posed, leading to questions of accuracy and stability of the results. Although Bastiaans illustrated that certain choices of window lead to particularly quirky representations and provided some insight into the mechanisms, much was left to be categorized. In the ten years following his valuable work, certain of the practical problems in performing the computations necessary to evaluate the transforms have come under careful scrutiny. Introduction of the Zak transform in the role of operational calculus for the Gabor theory has led to an improved understanding of the underlying mappings, and this in turn has shed light upon the supporting computational processes. Although this task is far from complete, much more is now known about the accuracy and stability of algorithms for computation of the Gabor transform [5, 6]. In addition, the computational complexity of these algorithms is also better understood [23, 24]. We can now safely pay a little more attention to what it is we ought to compute.

In deriving digital computation algorithms for the Gabor transforms, one must first make the equations discrete and finite. The most straightforward means is to truncate the time functions to some compact support and sample them at a rate that captures their significant behavior. This process relinquishes some control over the end product and blurs the

Nonetheless, some consistent theories can be developed in this manner. Wexler and Raz [27] have suggested a discrete formulation in which the time extent of the window function is truncated, but that of the biorthogonal function is not; furthermore, the frequency index remains unbounded. Those authors have commented that the unbounded supports make this form ill-suited to digital computation.

More recently, Balart [7] has developed a finite-dimension matrix formulation of the fully truncated case, in which it can be argued that the "true" discrete Gabor coefficients are produced when the window has bounded support. The approach of Auslander et. al. [2, 6], which uses the Zak transform, is fully discrete and finite, and produces an invertible transform when no sample of the Zak transform of the window function is zero. The coefficients produced by this algorithm, however, can differ considerably from those produced by Balart's method, as numerical experiments performed by Balart and the author have shown. Both methods are highly useful transforms that nevertheless differ slightly from the continuous parameter Gabor transform. Results of the two converge for windows having sharp discontinuities, deviating more and more from one another as the window appears to become continuous. This is not unexpected based on the Zero Theorem of Auslander and Tolimieri [4, Theorem II.2], which states that if the Zak transform of a sufficiently rapidly decaying function is continuous, it must have a zero. It is empirically true that the two methods are nearly the same if the Zak transform of the window has a narrow dynamic range, but they can diverge if the samples of that Zak transform make a close approach to a zero.

A discretization and finitization process in which some of the relations to the continuous parameter case are preserved would be more satisfying and would lend insight into how successful the more brute-force approach might prove. Wexler and Raz [27] have presented the sampled approach for periodic functions—one of the methods discussed in this paper—through a matrix formulation, and Auslander et. al. [2, 6] have given the Zak transform relationships under periodization of the signal. The contribution of the present paper is a unified treatment of discretization of the Gabor transform and its ancillary functions by

The Gabor discretization issue parallels what is done in transitioning from the Fourier integral to the finite DFT, where periodizing and sampling in the time domain likewise periodizes and samples the frequency domain picture. Moreover, the two periodic sequences so created are related by the finite discrete Fourier transform (DFT). The result for a function f and its Fourier transform F

$$F(v) = \int_{-\infty}^{\infty} dt f(t) \exp(-j2\pi vt), \quad (1)$$

is that the following two sums are a DFT pair [25]:

$$\left\{ \sum_k f(p + kP)T/P \right\} \Leftrightarrow \left\{ \frac{1}{T} \sum_j F(q + jP)/T \right\}; \quad p, q \in \mathbb{Z}. \quad (2)$$

For this rather remarkable formula—in which the essence of the Fourier transform property is maintained under sampling and aliasing—to hold, the sampling rates in the conjugate domains must be properly related. One samples f at spacing $t_s = T/P$, for positive T and positive integer P , and the resulting sequence is aliased at every P -th sample. The spectrum is then sampled at interval $f_s = 1/T$ and also aliased at the P -th sample, resulting in two sequences of period P . The relation

$$Pf_s t_s = 1 \quad (3)$$

characterizes the parameter constraint.

It is important to keep in mind that equation (2) tells us nothing about the resemblance of these finite discrete sequences to their parent, continuous parameter functions; however, conditions under which one is essentially a sampled version of the other are well known. In a clever work, Auslander and Grünbaum [3] have explored signal-independent error bounds for this comparison, in the process numerically verifying the necessity of condition (3).

~~Final TR: APPLICATION OF THE GABOR REPRESENTATION TO MILITARY PROBLEMS~~
We show that a relation corresponding to (2) holds for the Gabor transform. We then derive all the discrete finite Gabor transform equations from the continuous parameter versions through the method of periodization and sampling, in the process explicitly exhibiting the aliasings that permit one periodic sequence to be the finite, discrete Gabor transform of the other. By examining the various forms in which the Gabor equations can be expressed, we discover how the input, window, biorthogonal function, Gabor coefficients and Zak transforms map under periodization and sampling.

Section 2 briefly summarizes the continuous parameter Gabor equations, giving the fundamental definitions for and relationships among the: (i) signal, or function to be expanded, (ii) window function, (iii) biorthogonal function, (iv) Gabor coefficients, (v) Zak transforms, (vi) inner products in the time and Zak domains, and (vii) sampled auto- and cross-ambiguity functions of the signal and window. In Section 3, we transform the synthesis equation that recovers a function from its Gabor coefficients into discrete, periodic form. The inverse transformation is visited in Section 4 in terms of the biorthogonal function. In Section 5, the Gabor relations expressed in terms of Zak transforms are discretized. Behavior of the inner product of Zak transforms under the discretization process is shown in Section 6. Relationships that permit solution for the Gabor coefficients to be expressed as a deconvolution of the sampled crossambiguity function of the signal and window are analyzed under periodization and sampling in Section 7. In Section 8 we obtain the Gabor coefficients of the DFT of the given signal. Our concluding observations appear in Section 9.

2. The continuous parameter Gabor equations

In this section we summarize without proof results that serve as a basis for development of the discrete equations.

2.1. Definition (synthesis formula)

A Gabor representation of a time function $f(t)$ is a series expansion of the following form:

$$f(t) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{m,n} w(t - nT) \exp(j2\pi m t/T), \quad (4)$$

in which $\{a_{m,n}\}$, $m, n \in \mathbb{Z}$, are the *Gabor coefficients*, and $w(t)$, the *window*. Translates of $w(t)$ over the *von Neumann lattice* having a unit area cell of dimension $T \times (1/T)$ in the time-frequency plane form the *Gabor basis* $\{w_{m,n}(t)\}$, as follows:

$$w_{m,n}(t) = w(t - nT) \exp(j2\pi m t/T); m, n \in \mathbb{Z}. \quad (5)$$

In this regard, the Gabor expansion may be compared to its generalization, the Weyl-Heisenberg expansions, which can use any density of basis functions ≥ 1 [13, 14, 18].

In general, the $\{w_{m,n}(t)\}$ are nonorthogonal, and (4) is a nonorthogonal expansion, which is the reason one must take some pains to assure accuracy and stability of algorithms that compute the transformation. Gabor expansions can be orthogonal, but this is neither required nor always desirable, as discussed in [13]. Conditions for orthogonality of the Gabor basis are found in Boon and Zak [10]. These can be expressed in terms of the auto-ambiguity function of the window—see Tolimieri and Orr [26]—and a general construction of orthogonal Gabor expansions has recently been developed by Coifman *et. al* [11].

which we denote by $g(t)$, and define in notation different from Gabor's:

$$g(t) = \left(\frac{1}{\sigma}\right)^{1/2} \exp\left[-\frac{\pi}{2} \left(\frac{t}{\sigma}\right)^2\right], \quad (6)$$

where σ is the RMS pulse width. The normalization in (6) is that the window has unit energy, or L_2 norm, which is a convenient condition we shall assume throughout for all windows:

$$\|g\|^2 = \int_{-\infty}^{\infty} dt |g(t)|^2 = 1. \quad (7)$$

Helstrom [19], in 1966, demonstrated the continuous parameter version of Gabor's Gaussian expansion and noted the relationship to Glauber's coherent states of quantum mechanics [16]. In 1967 both Montgomery and Reed [22] and Crum [12] published on the expansion of functions in other than Gaussian elementary signals, generating what was then viewed as a continuous parameter Gabor-like transform that would later be called the short-time Fourier transform.

It was in Bastiaans [8] that the biorthogonal function for the Gaussian window was introduced. In later considering other functions for the window role, Bastiaans [9] provided the basis of modern Gabor theory by his concern for the inversion of Gabor's representation.

2.2. The biorthogonal function (analysis formula)

In the *biorthogonal* method, one uses both the basis functions $\{w_{m,n}(t)\}$ and a related set of *biorthogonal functions* $\{b_{m,n}(t)\}$,

$$b_{m,n}(t) = b(t - nT) \exp(j2\pi m t/T); \quad m, n \in \mathbb{Z}. \quad (8)$$

$$(w_{m,n} | b_{p,q}) = \int_{-\infty}^{\infty} dt w(t - nT) b^*(t - qT) \exp(j2\pi(m - p)t/T) = \delta_{n,q} \delta_{m,p}. \quad (9)$$

When such $b(t)$ exists, (9) permits analysis of f into its coefficients by inner products of $f(t)$ and the $\{b_{m,n}(t)\}$ according to the formula

$$a_{m,n} = (f | b_{m,n}) = \int_{-\infty}^{\infty} dt f(t) b_{m,n}^*(t). \quad (10)$$

2.3. The Zak transform

A second method of writing the Gabor coefficients is based upon using the *Zak transform*, a time-frequency mapping given by

$$Zf(v, \tau) = \sum_{k=-\infty}^{\infty} f(kT + \tau) \exp(-j2\pi k\tau/T). \quad (11)$$

Taking Zak transforms on both sides of (4) yields the following relation among the Zak transform of f , the Zak transform of w , and the Gabor coefficients [2, 6, 21]:

$$Zf(v, \tau) = Zw(v, \tau) \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{m,n} \exp[j2\pi(m\tau/T - n\tau/T)]. \quad (12)$$

The $\{a_{m,n}\}$ are found formally by inverting the two-dimensional Fourier series on the right:

$$a_{m,n} = \int_0^T d\tau \int_0^{1/T} dv \frac{Zf(v, \tau)}{Zw(v, \tau)} \exp(j2\pi(-m\tau/T + n\tau/T)). \quad (13)$$

If, however, Zw has a zero where Zf does not, Zf/Zw is not an $L^2(T \times F)$ function, where $T = [0, T)$ and $F = [0, 1/T)$. Zf/Zw can be $L^1(T \times F)$, in which case Gabor coefficients can be computed from (13), but may not be square summable, and any truncated representation that uses but a finite number of them can have large errors. Also, a Riemann sum approximation to the double integral may not provide numerically accurate formulas. Issues related to stabilization of this process are not discussed in this paper.

The relationship between the window and biorthogonal functions is reciprocal in the Zak domain: $Zb(v, \tau) = 1/[TZ^*w(v, \tau)]$.

Synthesis of f from (11) follows by inverting the DFT after recovering $Zf(v, \tau)$ from the $\{a_{m,n}\}$ via (12):

$$f(kT + \tau) = T \int_0^{1/T} dv Zf(v, \tau) \exp(j2\pi kv\tau), \quad k \in \mathbb{Z}, \tau \in T. \quad (14)$$

2.4. The deconvolution methods

Deconvolution permits recovery of f from its sampled short-time Fourier transform (STFT) $\{(f|w_{m,n})\}$, using the fact that the Zak transform is an isomorphism, preserving inner products to within a scale factor [21]:

$$(Zf|Zg) = \frac{1}{T} (f|g). \quad (15)$$

where

$$\langle Zf | Zg \rangle = \int_0^{1/T} dv \int_0^T d\tau Zf(v, \tau) Z^*g(v, \tau). \quad (16)$$

In [18] and [21], the authors define the Zak transform with a normalization that makes the scale factor in (15) unity for all T . Multiplying both sides of (12) by $Z^*w(v, \tau)$,

$$Zf(v, \tau) Z^*w(v, \tau) = |Zw(v, \tau)|^2 \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{m,n} \exp[j2\pi(mt/T - n\tau T)], \quad (17)$$

and computing the Fourier coefficients of both expressions reduces (17), after simplification, to the convolution [2]

$$\langle f | w_{m,n} \rangle = \sum_{p=-\infty}^{\infty} \sum_{q=-\infty}^{\infty} a_{p,q} \langle w | w_{m-p, n-q} \rangle \quad (18)$$

This equation can of course be obtained more directly from (4) by taking the inner product indicated on the left and employing the shift theorem [20]. For present purposes it is the Fourier transform of (17) that is of greater interest.

Equation (18) is often used in comparing the Gabor and short-time Fourier transforms [1]. It shows the 'smearing' of the Gabor coefficients by the window-dependent kernel to produce the STFT. If a function w is used as a window in both transforms it cannot have both the smoothness desired for spectral analysis and the frame structure that stabilizes the Gabor expansion, as explained by the Balian-Low theorem [17].

Extraction of the Gabor coefficients from (18) requires deconvolution of the sampled STFT. Fourier analysis of both sides of (18) converts (17) into:

$$\{a_{m,n}\} = \text{DFT}^{-1} \left(\frac{\sum_{r=-\infty}^{\infty} \sum_{s=-\infty}^{\infty} (f | w_{r,s}) \exp [j2\pi(rvT - svT)]}{|Zw(v, \eta)|^2} \right) \quad (19)$$

Deconvolution is also the generalized inverse for incomplete Gabor expansions, providing the coefficients of the orthogonal projection of f onto the subspace spanned by the basis vectors.

In the following sections we convert the formulas of each method to a self-consistent set of equations that describe the finite, discrete Gabor transform, obtaining tools that are amenable to digital computation of the transforms.

3. The synthesis formula

In this and the following sections we develop the finite discrete Gabor relationships. Derivation of the first formula is presented in full. Following that derivations are relegated to the Appendix.

The starting point is (4), from which we eventually want to prove (33), which expresses the samples of the periodized signal (29) as a function of suitably periodized window samples (30) and Gabor coefficients (32). We periodize $f(t)$ to a period that we restrict to a multiple of T , NT :

$$f_{NT}(t) = \sum_{k=-\infty}^{\infty} f(t + kNT), \quad (20)$$

and formally write the Gabor expansion of $f_{NT}(t)$ as

$$f_{NT}(t) = \sum_{k=-\infty}^{\infty} \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{m,n} w(t - nT + kNT) \exp [j2\pi m(t - kNT)/T]. \quad (21)$$

If we make the following rearrangement of (21),

$$f_{NT}(t) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{m,n} \left[\sum_{k=-\infty}^{\infty} w[t - (n - kN)T] \right] \exp[j2\pi mt/T], \quad (22)$$

we see that the periodization of f results in a corresponding periodization of window, $w_{NT}(t)$,

$$w_{NT}(t) = \sum_{k=-\infty}^{\infty} w(t + kNT), \quad (23)$$

yielding

$$f_{NT}(t) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} a_{m,n} w_{NT}(t - nT) \exp[j2\pi mt/T]. \quad (24)$$

By (24), the Gabor coefficients are invariant to this joint periodization of the signal and window.

An interpretation of more immediate use results from a rearrangement of (24). Replace n by its modulo N representation, $n = q + sN$, where $0 \leq q \leq N - 1$, $s \in \mathbb{Z}$, and replace the corresponding sum by a double sum on q and s :

$$f_{NT}(t) = \sum_{m=-\infty}^{\infty} \sum_{s=-\infty}^{\infty} \sum_{q=0}^{N-1} a_{m,q+sN} w_{NT}(t - qT) \exp[j2\pi mt/T]. \quad (25)$$

The periodicity of w has been exploited in reaching (25). Reorder the summations so that the innermost is over s , the time index of the coefficients:

$$f_{NT}(t) = \sum_{m=-\infty}^{\infty} \sum_{q=0}^{N-1} \left[\sum_{s=-\infty}^{\infty} a_{m,q+sN} w_{NT}(t - qT) \exp(j2\pi ms/T) \right] \quad (26)$$

The bracketed term in (26) is a time-aliased Gabor coefficient

$$a_{m,q}^{(N)} = \sum_{s=-\infty}^{\infty} a_{m,q+sN} \quad 0 \leq q \leq N-1. \quad (27)$$

By (26) and (27), periodizing $f(t)$ similarly periodizes the window and the $\{a_{m,n}\}$ in time.

Now sample $f_{NT}(t)$ at the M -th harmonic of $1/T$, so that the k -th sample occurs at time kT/M . The resulting sampled sequence then has period $P = MN$, and (26) becomes:

$$f_k^{(P)} = \sum_{m=-\infty}^{\infty} \sum_{q=0}^{N-1} a_{m,q}^{(N)} w_k^{(P)} \exp(j2\pi mk/M), \quad (28)$$

where we have introduced simpler notation for the periodized data and window:

$$f_k^{(P)} = f_{NT}(kT/M) \quad (29)$$

$$w_k^{(P)} = w_{NT}(kT/M). \quad (30)$$

Now apply modular representation to the frequency index m , since the complex exponential in (28) has period M in m . Let $m = p + rM$, $0 \leq p \leq M-1$, $r \in \mathbb{Z}$, and reorder the sums:

$$f_k^{(P)} = \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} \left[\sum_{r=-\infty}^{\infty} a_{p+rM,q}^{(N)} \right] w_k^{(P)} \exp[j2\pi k(p + rM)/M]. \quad (31)$$

This operation periodizes the Gabor coefficients in frequency, motivating the definition

$$a_{p,q}^{(M,N)} = \sum_{r=-\infty}^{\infty} a_{p+rM,q}^{(N)} = \sum_{r=-\infty}^{\infty} \sum_{s=-\infty}^{\infty} a_{p+rM,q+sN}. \quad (32)$$

Combining all the above we have

$$f_k^{(P)} = \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} w_k^{(P)} \exp(j2\pi kp/M). \quad (33)$$

In summary, the finite discrete Gabor representation of the periodized and sampled f uses the periodized and sampled window function and the doubly periodized Gabor coefficients of the parent continuous parameter signal.

4. The biorthogonal formula

This method begins with (10), which expresses the Gabor coefficients as inner products of f and the set of biorthogonal functions, and executes the periodization (32). Following steps similar to the above we find

$$a_{m,n}^{(M,N)} = \frac{T}{M} \sum_{r=0}^{P-1} f_r^{(P)} b_{r-mN}^{(P)} \exp(j2\pi nr/M). \quad (34)$$

Thus the process of sampling and periodizing replaces the integral expression for $a_{m,n}$ by a finite, discrete inner product of the sampled, periodized data and a similarly sampled and periodized set of biorthogonal functions. The factor T/M in (34) is the sampling interval and takes on the role played by the differential dt in (10).

PART II: APPLICATION OF THE GABOR REPRESENTATION TO MILITARY PROBLEMS

5. The Zak Transform

In (12), which states the relationship of the Zak transforms of the signal and window, we can replace the index variables as follows:

$$\begin{aligned} m &= p + rM; 0 \leq p \leq M-1, r \in \mathbb{Z} \\ n &= q + sN; 0 \leq q \leq N-1, s \in \mathbb{Z} \end{aligned} \quad (35)$$

and sample in time and frequency, i.e.

$$\nu = \frac{n}{NT}, \quad \tau = \frac{mT}{M} \quad (36)$$

to get

$$(Zf)_{n,m} = (Zw)_{n,m} \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \exp(j2\pi(mp/M - nq/N)). \quad (37)$$

Above we have introduced a shorthand notation for the Zak transforms,

$$Zf(n/NT, mT/M) = (Zf)_{n,m}. \quad (38)$$

The intermediate steps are found in the Appendix. Inversion of the 2-D Fourier series immediately yields

$$a_{m,n}^{(M,N)} = \frac{1}{MN} \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} \frac{(Zf)_{n,p}}{(Zw)_{n,p}} \exp[-j2\pi(mp/M - nq/N)]. \quad (39)$$

Next we show how to recover the finite discrete version of a function from its Zak transform, starting from (14) and periodizing. The periodization effects sampling in the frequency variable at spacing $1/NT$. Then sampling at $\tau = pT/M$, $0 \leq p \leq M-1$, yields

$$f_{NT}\left[\left(\frac{p}{M} + k\right)T\right] = f_{p+kM}^{(P)} = \frac{1}{N} \sum_{q=0}^{N-1} (Zf)_{q,p} \exp(j2\pi kq/N); \quad 0 < p < M-1, k \in \mathbb{Z} \quad (40)$$

Formulas (37) and (39) show that the discrete Gabor coefficients will exist either for all f in $L^2(\mathbb{Z}/P)$ when Zw has no zero on the sampling grid, or in general, in the subspace of functions f whose Zak transforms vanish at the zeros of Zw . Via (40), the P samples of f can be recovered from the P samples of its Zak transform by a set of DFT's. Because of its quasi-periodicity, the Zak transform acquires no additional periodicities in the process.

6. The inner product formulas

To streamline the derivation of the finite discrete deconvolution equations in Section 7, we will need to know how inner products map under periodization and sampling. Let $f^{(P)}$ and $g^{(P)}$ be two periodized, sampled functions. A natural definition of their inner product is

$$\langle f^{(P)} | g^{(P)} \rangle = \sum_{p=0}^{P-1} f_{NT}\left(\frac{pT}{M}\right) g_{NT}^*\left(\frac{pT}{M}\right) = \sum_{p=0}^{P-1} f_p^{(P)} (g_p^{(P)})^*. \quad (41)$$

We find that the result may be written in either of the forms

$$\begin{aligned} \langle f^{(P)} | g^{(P)} \rangle &= \frac{1}{N} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (Zf)_{n,m} (Zg)_{n,m}^* \\ &= \frac{1}{N} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (Zf)_{n,m} (Zg^*)_{N-n,m}. \end{aligned} \quad (42)$$

Thus the inner product of two periodized, sampled functions is proportional to the 2-D inner product of their Zak transforms.

7. The ~~STANDARD~~ APPLICATION OF THE GABOR REPRESENTATION TO MILITARY PROBLEMS

The equation for the deconvolution of the periodized, sampled Gabor coefficients from the sampled short-time Fourier transform is readily found from (33). Take the indicated inner product and apply the shift theorem to the indices, yielding

$$\langle f^{(P)} | w_{r,s}^{(P)} \rangle = \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \langle w^{(P)} | w^{(P)}_{r-p,s-q} \rangle. \quad (43)$$

In the appendix we show how this result may be found via the Zak transform.

The correspondence of (43) to (18) is clear. The doubly infinite, two-dimensional convolution has been replaced by a finite, end-around convolution. Again, the doubly aliased Gabor coefficients take the role formerly played by the Gabor coefficients of $f(r)$.

8. Gabor transform of the DFT

The Gabor representation of a signal and its Fourier transform are tightly interrelated by way of the Fourier transform of the window. If we begin with the discrete representation (33) and take the DFT, we have

$$\hat{f}_r^{(P)} = \sum_{k=0}^{N-1} f_k^{(P)} \exp(-j2\pi kr/P) = \sum_{k=0}^{N-1} \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} w_{k-q}^{(P)} \exp[j2\pi(mN - r)k/P]. \quad (44)$$

Performing the summation on k first isolates the expression

$$\hat{w}_r^{(P)} = \sum_{k=0}^{N-1} w_k^{(P)} \exp(-j2\pi ks/P), \quad (45)$$

which is the DFT of the window. In terms of this quantity, the transform becomes

$$\hat{f}_r^{(P)} = \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \hat{w}_{r-mN}^{(P)} \exp[-j2\pi(r-mN)q/N]. \quad (46)$$

In (46), the r -th DFT coefficient is expressed via a Gabor expansion whose coefficients are those of f and using window \hat{w} . The roles of time and frequency have been interchanged and there is a sign change in the exponential. This formula can be put in an alternative form that recovers some of the lost symmetry by replacing q with $N - n$ and denoting the Gabor coefficients of \hat{f} as

$$\hat{a}_{n,m}^{(N,M)} = a_{m,N-n}^{(M,N)}. \quad (47)$$

Exchange of time and frequency roles is now complete; the coefficients of \hat{f} with respect to \hat{w} are now the coefficients of f rotated 90° in the time-frequency plane:

$$\hat{f}_r^{(P)} = \sum_{n=0}^{N-1} \sum_{m=0}^{M-1} \hat{a}_{n,m}^{(N,M)} \hat{w}_{r-mN}^{(P)} \exp[j2\pi(r-mN)n/N]. \quad (48)$$

9. Observations

The equations defining the continuous parameter Gabor transform have been converted to finite, discrete form using periodization and sampling. The periodic sequences $\{f_k^{(P)}\}$ and $\{a_{p,q}^{(M,N)}\}$, where $0 \leq k \leq P - 1$, $0 \leq p \leq M - 1$, $0 \leq q \leq N - 1$, become a Gabor transform pair. In the linear transformation relating these, the window function is replaced by the periodic sequence $\{w_k^{(P)}\}$, $0 \leq k \leq P - 1$. When the biorthogonal function is well behaved, the $\{a_{p,q}^{(M,N)}\}$ can be obtained from an inner product of the data, $\{f_k^{(P)}\}$, and a sampled, periodized version of b . Relationships involving the Zak transforms map similarly; the continuous arguments are replaced by discrete ones, $Zf \rightarrow \{(Zf)_{q,p}\}$, without periodization. Likewise, the convolution equation for the Gabor coefficients converts to a

In applying the Gabor transform to signal processing problems, one would like to have a clear correspondence between the finite, discrete formulation used in computation and the continuous parameter equations of the theory. Specifically, one wants to see the analog signal, window and biorthogonal functions map into a finite set of their samples, and to have a finite set of Gabor coefficients that (i) are defined over a time-frequency region where the coefficients of the analog signal have significant values, and (ii) are approximately equal to the latter coefficients in that region. Creating the correspondence by periodization and sampling makes this relatively straightforward. Since Gabor coefficients are aliased in time and frequency in the discretization, one needs to assure that the replicas are sufficiently spread to prevent significant overlap. This is done by choosing the signal period to be suitably long with respect to the lengths of the signal and window functions, and by sampling at a rate high enough—essentially a composite Nyquist rate—to capture the significant frequency behavior of the signal and window. A sampling theory for Gabor expansions will be the subject of a forthcoming paper by the author, but in its absence one can still be confident that when these guideline precautions are observed, computations made with the finite, discrete Gabor transform should faithfully reflect the behavior of their continuous parameter counterpart.

References

1. L. Auslander, C. Buffalano, R. Orr and R. Tolimieri, "A Comparison of the Gabor and Short-Time Fourier Transforms for Signal Detection and Feature Extraction in Noisy Environments," Proc. SPIE, Vol. 1348, November 1990, pps. 230-247.
2. L. Auslander, I. Gertner and R. Tolimieri, "The Discrete Zak Transform Application to Time-Frequency Analysis and Synthesis of Non-Stationary Signals," IEEE Trans. SP, Vol. 39, No. 4, April 1991, pps. 825-835.
3. L. Auslander and A. Grünbaum, "The Fourier transform and the discrete Fourier

Atlantic Aerospace Electronics Corporation
transform." Inverse Problems, Vol. 5, 1989, pp. 149-164.

4. L. Auslander and R. Tolimieri, Abelian Harmonic Analysis. Theta Functions and Function Algebras on a Nilmanifold, Springer-Verlag, New York, 1975, Ch. II, p. 18.
5. L. Auslander and R. Tolimieri, "Computing Decimated Finite Cross-Ambiguity Functions" IEEE Trans. ASSP, Vol. 36, No. 3, March 1988, pp. 359-363.
6. L. Auslander and R. Tolimieri, "On finite Gabor expansions of signals," Signal Processing. Part I: Signal Processing Theory, Springer-Verlag, New York, 1990, pp. 13-23.
7. R. Balart, "Matrix reformulation of the Gabor transform," Optical Engineering, Vol. 31, No. 6, pp. 1235-1242, June, 1992.
8. M. J. Bastiaans, "Gabor's Expansion of a Signal into Gaussian Elementary Signals," Proc. IEEE, Vol. 68, No. 4, April 1980, pp. 538-539.
9. M. J. Bastiaans, "A Sampling Theorem for the Complex Spectrogram, and Gabor's Expansion of a Signal in Gaussian Elementary Signals," Optical Engineering, Vol. 20, No. 4, July/August 1981, pps. 594-598.
10. M. Boon and J. Zak, "Amplitudes on von Neumann lattices," J. Math. Phys., Vol. 22, pp. 1090-1099, 1981.
11. R. Coifman, Y. Meyer and V. Wickerhauser, Wavelet Analysis and Signal Processing, Yale U., 1991.
12. M. M. Crum, "Representation of Signals," IEEE Trans. IT, April 1967, p. 348.
13. I. Daubechies, A. Grossmann and Y. Meyer, "Painless nonorthogonal expansions," J. Math. Phys., Vol. 27, No. 5, May 1986, pp. 1271-1283.
14. I. Daubechies, "The Wavelet Transform, Time-Frequency Localization and Signal Analysis," IEEE Trans. IT, Vol. 36, No. 5, pp. 961-1005, September, 1990.
15. D. Gabor, "Theory of Communication," IIEE, Vol. 93, 1946, pps. 429-459.
16. R. J. Glauber, "Coherent and incoherent states of the radiation field," Phys. Rev., Vol. 131, September 1963, pp. 2766-2788.
17. C. Heil, Uncertainty Principles and the Balian-Low Theorem, preprint, The MITRE Corp., 1990.

- ~~Final TR: APPLICATION OF THE GABOR REPRESENTATION TO MILITARY PROBLEMS~~
18. C. Heil and D. Walnut, "Continuous and Discrete Wavelet Transforms." SIAM REVIEW
Vol. 31, No. 4, 1989, pp. 628-666.
 19. C. W. Helstrom, "An Expansion of a Signal in Gaussian Elementary Signals." IEEE Trans. IT, 1966, pp. 81-82.
 20. A. J. E. M. Janssen, "Gabor Representation of Generalized Functions." J. Math. Anal. and Appl., Vol. 83, pp. 377-394, 1981.
 21. A. J. M. Janssen, "The Zak Transform: A Signal Transform for Sampled Time-Continuous Signals." Phillips J. Res., Vol. 43, 1988, pps. 23-69.
 22. L. K. Montgomery, and I. S. Reed, "A Generalization of the Gabor-Helstrom Transform." IEEE Trans. IT, April 1976, pp. 234-235.
 23. R. S. Orr, "Computational Assessment of Gabor Representations." Proc IEEE ICASSP'91, Vol. 3, Digital Signal Processing, May 1991, pp. 2217-2220.
 24. R. S. Orr, "The Order of Computation for Finite Discrete Gabor Transforms." to appear. IEEE Trans. SP, Vol. 41, No. 1, January 1993, pp. 122-130.
 25. W. Rudin, Real and Complex Analysis, McGraw-Hill Book Company, New York, NY, 1974, p. 210.
 26. R. Tolimieri, and R. S. Orr, "Poisson Summation, the Ambiguity Function and the Theory of Weyl-Heisenberg Frames." accepted for publication pending revision, IEEE Trans. IT, submitted July 1990.
 27. J. Wexler and S. Raz, "Discrete Gabor Expansions." Signal Processing, No. 21, 1990, pp. 207-220.

Appendix

Proofs for Section 4:

Writing the time index in (10) in mod N representation gives

$$a_{m,n+qN} = \int_{-\infty}^{\infty} dt f(t + qNT) b_{m,n}^*(t), \quad (\text{A1})$$

and therefore

$$\sum_{q=-\infty}^{\infty} a_{m,n+qN} = \int_{-\infty}^{\infty} dt \left[\sum_{q=-\infty}^{\infty} f(t + qNT) \right] b_{m,n}^*(t). \quad (\text{A2})$$

Periodizing the Gabor coefficients with respect to frequency yields

$$a_{m,n}^{(M,N)} = \sum_{q=-\infty}^{\infty} \sum_{p=-\infty}^{\infty} a_{m+pM,n+qN} = \int_{-\infty}^{\infty} dt \left[\sum_{q=-\infty}^{\infty} f(t + qNT) \right] \left[\sum_{p=-\infty}^{\infty} b_{m+pM,n}^*(t) \right]. \quad (\text{A3})$$

The biorthogonal summation in (A3) can be further developed:

$$\sum_{p=-\infty}^{\infty} b_{m+pM,n}^*(t) = \sum_{p=-\infty}^{\infty} b_{m,n}^*(t) \exp(-j2\pi pMt/T) = b_{m,n}^*(t) \sum_{p=-\infty}^{\infty} \exp(-j2\pi pMt/T). \quad (\text{A4})$$

The summation of complex exponentials can be recognized as proportional to the formal Fourier series of a train of impulses separated by T/M . The r -th Fourier coefficient of such a train of unit impulses is given by

$$\frac{M}{T} \int_0^{TM} dt \left[\sum_{k=-\infty}^{\infty} \delta(t - kTM) \right] \exp(-j2\pi rMt/T) = \frac{M}{T}. \quad (A5)$$

implying

$$\sum_{p=-\infty}^{\infty} b_{m+PM,n}^*(t) = b_{m,n}^*(t) \left(\frac{T}{M} \right) \sum_{k=-\infty}^{\infty} \delta(t - kTM). \quad (A6)$$

When this is inserted into (A3), the impulse functions sample the biorthogonal functions, yielding

$$a_{m,n}^{(M,N)} = \frac{T}{M} \sum_{k=-\infty}^{\infty} f_k^{(P)} b_{m,n}^*\left(\frac{kT}{M}\right), \quad (A7)$$

where $f_k^{(P)}$ is the sampled, periodized f , see definition (29). We exploit the periodicity by writing the index k in terms of its quotient and remainder modulo P , $k = r + Ps$, rearranging the sum to yield

$$\begin{aligned} a_{m,n}^{(M,N)} &= \frac{T}{M} \sum_{r=0}^{P-1} \sum_{s=-\infty}^{\infty} f_{r+Ps}^{(P)} b_{m,n}^*\left(\frac{(r+Ps)T}{M}\right) \\ &= \frac{T}{M} \sum_{r=0}^{P-1} f_r^{(P)} \left[\sum_{s=-\infty}^{\infty} b_{m,n}^*\left(\frac{(r+Ps)T}{M}\right) \right]. \end{aligned} \quad (A8)$$

The bracketed term is the periodized, sampled biorthogonal function that can be written in the notation of (29) and (30) as

$$b_{r-mN}^{(P)} \exp(j2\pi mr/M) = \sum_{s=-\infty}^{\infty} b_{m,m} \left(\frac{(r + sP)T}{M} \right) \quad (A9)$$

letting us write the final result as in (34).

Proofs for Section 5:

Beginning with (12), and making the change of variables indicated in (35), we have

$$Zf(v, \tau) = ZW(v, \tau) \sum_{r=-\infty}^{\infty} \sum_{s=-\infty}^{\infty} \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p+rM, q+sN} \exp\{j2\pi[(p+rM)\tau/T - (q+sN)vT]\} \quad (A10)$$

Sampling (A10) on the grid given in (36) yields

$$(Zf)_{n,m} = (ZW)_{n,m} \sum_{r=-\infty}^{\infty} \sum_{s=-\infty}^{\infty} \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p+rM, q+sN} \exp\{j2\pi[(p+rM)m/M - (q+sN)n/N]\}. \quad (A11)$$

Upon simplifying the exponential terms and recognizing the inner double sum of the Gabor coefficients over p and q as the doubly periodic coefficients, (A11) becomes

$$\begin{aligned} f_{NT}(\tau + kT) &= \sum_{q=-\infty}^{\infty} f(\tau + kT + NqT) \\ &= \sum_{q=-\infty}^{\infty} T \int_0^{1/T} dv Zf(v, \tau) \exp[j2\pi(k - Nq)vT] \\ &= T \int_0^{1/T} dv Zf(v, \tau) \exp[j2\pi k vT] \sum_{q=-\infty}^{\infty} \exp[j2\pi Nq vT]. \quad (A12) \end{aligned}$$

$$(Zf)_{n,m} (Zw_{r,s})_{n,m}^* = (Zw)_{n,m} (Zw_{r,s})_{n,m}^* \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \exp[j2\pi(mp/M - nq/N)]. \quad (A21)$$

Now sum both sides of (A21) over m and n . The left and right sides become, respectively,

$$\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (Zf)_{n,m} (Zw_{r,s})_{n,m}^* = N \langle f^{(P)} | w_{r,s}^{(P)} \rangle \quad (A22)$$

and

$$\begin{aligned} & \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (Zw)_{n,m} (Zw_{r,s})_{n,m}^* \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \exp[j2\pi(mp/M - nq/N)] \\ &= \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (Zw)_{n,m} (Zw_{r,s})_{n,m}^* \exp[j2\pi(mp/M - nq/N)] \\ &= \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} (Zw)_{n,m} (Zw_{r-p,s-q})_{n,m}^* \\ &= N \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \langle w^{(P)} | w^{(P)}_{r-p,s-q} \rangle \end{aligned} \quad (A23)$$

Combining the above, we have (43).

$$\begin{aligned}
 f_{NT}(\tau + kT) &= T \int_0^{1/T} dv Zf(v, \tau) \exp(j2\pi k v T) \sum_{q=-\infty}^{\infty} \frac{1}{NT} \delta(v - \frac{q}{NT}) \\
 &= \frac{1}{N} \sum_{q=0}^{N-1} Z\left(\frac{q}{NT}, \tau\right) \exp(j2\pi k q / N).
 \end{aligned} \tag{A13}$$

Proofs for Section 6:

We derive the inner product formula via the Zak transform in this section. Equation (41) can be rewritten replacing p with its representation mod M , $p = nM + m$, where $0 \leq n \leq N - 1$, and $0 \leq m \leq M - 1$:

$$\langle f^{(P)} | g^{(P)} \rangle = \sum_{m=0}^{M-1} \sum_{r=0}^{N-1} f_{m+rM}^{(P)} (g_{m+rM}^{(P)})^* \tag{A14}$$

Now replace the samples $f_p^{(P)}$ and $g_p^{(P)}$ by their discrete Zak transform representations (40) to get

$$\begin{aligned}
 \langle f^{(P)} | g^{(P)} \rangle &= \sum_{m=0}^{M-1} \sum_{r=0}^{N-1} \left[\frac{1}{N} \sum_{s=0}^{N-1} (Zf)_{s,m} \exp(j2\pi sr/N) \right] \left[\frac{1}{N} \sum_{n=0}^{N-1} (Zg)_{n,m} \exp(j2\pi nr/N) \right]^* \\
 &= \frac{1}{N^2} \sum_{m=0}^{M-1} \sum_{s=0}^{N-1} \sum_{n=0}^{N-1} (Zf)_{s,m} (Zg)_{n,m}^* \sum_{r=0}^{N-1} \exp[j2\pi(s-n)r/N].
 \end{aligned} \tag{A15}$$

Now recognize that

$$\sum_{r=0}^{N-1} \exp[j2\pi(s-r)N] = N\delta_{s-r} \quad (\text{A16})$$

so that (A15) takes the form given in (41).

Using the easily proved relationship [20]

$$Zf^*(v, \tau) = Z^*f(-v, \tau) \quad (\text{A17})$$

and the periodicity of the Zak transform in its frequency variable. The Zak transform of g can be expressed in terms of the conjugate of g :

$$(Zg)_{n,m}^* = (Zg^*)_{-n,m} = (Zg^*)_{N-n,m} \quad (\text{A18})$$

from which we get (42).

Proofs for Section 7:

We multiply both sides of (37) by $(Zw)_{n,m}^*$, finding

$$(Zf)_{n,m} (Zw)_{n,m}^* = |(Zw)_{n,p}|^2 \sum_{p=0}^{M-1} \sum_{q=0}^{N-1} a_{p,q}^{(M,N)} \exp[j2\pi(mp/M - nq/N)]. \quad (\text{A19})$$

If we multiply (A19) by $\exp[j2\pi(rp/M - sq/N)]$, the left-hand side becomes

$$(Zf)_{n,m} (Zw)_{n,m}^* \exp[j2\pi(rp/M - sq/N)] = (Zf)_{n,m} (Zw_{r,s})_{n,m}^* \quad (\text{A20})$$

where the second factor on the right is interpreted as the (n, m) -th sample of the Zak transform of $w_{r,s}$. The right-hand side of (A19) transforms similarly, leading to

APPENDIX D
MINIMUM DIMENSION GABOR PAPER

EXPERIMENTS IN DIMENSIONALLY-OPTIMUM GABOR REPRESENTATIONS

J. Sweeney, R. Orr, R. Balart, D. Buchanan and A. Humen

Atlantic Aerospace Electronics Corporation, 6404 Ivy Lane, Suite 300, Greenbelt, MD 20770 USA

ABSTRACT

Series expansions of signals in which significant features of the signal are captured in a few large coefficients are desirable. This work shows that given a collection of signals, it is possible to find Gabor representations for these that are maximally concentrated in time-frequency space. The problem addressed is: given a signal set, find the window function of the Gabor expansion that minimizes an "average dimension" of the signal representations relative to that window. The dimension measure employed is entropy based and related to the quantum-mechanical technique where one interprets expansion coefficients as probabilities.

An iterative algorithm based on partial derivatives of the signal set dimension with respect to the expansion function was used to evaluate effectiveness of several nonlinear optimization algorithms in finding an optimum window.

1. DIMENSIONALLY OPTIMUM REPRESENTATION

Series representations of signals in which most significant features of the signal are captured in a few large coefficients are desirable in problems of detection, feature extraction, characterization and data compression. Many standard expansions—e.g., classical Fourier methods—have but little flexibility to adapt to characteristics of the data. Time-frequency or time-scale representations encompassing a family of transforms specified by a window function (or analyzing wavelet), however, have great potential to accommodate the data under analysis by selection of a well-matched transform. To date there seems to have been little effort expended in achieving some of this potential for economy of representation.

A theory for selection of good basis functions is described in a prior publication [1]. This theory centers on defining a *dimension*—first for a signal and then for a

signal set—in terms of the expansion coefficients generated by the signals under a particular representation. In this paper we report early experiments in implementing optimization techniques to use the theory in obtaining dimensionally-optimum Gabor expansions.

2. THE DIMENSION OF SIGNALS

Time-bandwidth product is a familiar measure of the complexity of a signal. The *BT* product is intrinsic, i.e., independent of the representation applied. By contrast, the measure sought here is to be used in preferentially selecting among representations, and thus should be representation dependent. The same should be true of the dimension of a set of signals, although such a measure begins to take on a more intrinsic nature as suitable constraints are placed upon the eligible representations.

Since dimension is tied to the representation in force, all its properties should be expressible through the signal's expansion coefficients. When the expansion is a Gabor representation of $s_i(t)$,

$$s_i(t) = \sum_m \sum_n a_{m,n}^i w_{m,n}(t), \quad (1)$$

where the $\{w_{m,n}(t)\}$ are time and frequency translations of a window function $w(t)$ [2], dimension is expressed in terms of the $\{a_{m,n}^i\}$.

We have imposed six 'reasonableness' requirements on a definition of dimension, each expressible in terms of the signal, the expansion functions, and the coefficients—see [1] for details. Under these requirements the unique solution for the dimension has the form

$$D_i(s) = \exp \left[- \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} p_{m,n}^i \ln p_{m,n}^i \right], \quad (2)$$

$$p_{m,n}^i = \frac{|a_{m,n}^i|^r}{\sum_{m=0}^{M-1} \sum_{n=0}^{N-1} |a_{m,n}^i|^r} : r \geq 1. \quad (3)$$

This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by the U.S. Army Missile Command under contract No. DAAH01-91-C-R271. The United States Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright notation hereon.

The result is an entropy-derived measure for which probabilities are developed from the coefficients $\{a_{m,n}\}$, by analogy to quantum mechanical methods. When a set of signals $\{s_i(t)\}$ is considered, a *set dimension* can be defined as a weighted average of individual dimensions:

$$D = \sum \mu_i D_i \quad (4)$$

where the weights $\{\mu_i\}$ could represent probabilities, or could be assigned to reflect relative importance of the various signals. This same dimension measure has been proposed independently by Coifman *et al.* [3].

3. SOLUTION TECHNIQUE

Analytic minimization of (4) is in general a bleak prospect. Even writing the gradient of D with respect to the window function is difficult, and the solution of $\nabla D = 0$ is not apparent. We address the minimization as follows: (1) since expansion coefficients are linearly related to the biorthogonal function, we solve not for the window function w that minimizes D , but its biorthogonal, b —permitting the analytic gradient to be obtained—from which w is found; (2) solution for the optimum b is accomplished by iterative methods of nonlinear optimization theory, including both those that directly seek the minimum of D and those that solve for a root of the gradient of D . Derivative aiding can range from none to analytically provided second derivatives (Hessian matrix) [4]. Routines from the MATLAB Optimization Toolbox and NAG FORTRAN Library generated the results.

The equations resulting from the above are

$$\frac{\partial D}{\partial b_q} = D \frac{\partial \ln D}{\partial b_q} = 0; \quad 0 \leq q \leq P-1. \quad (5)$$

To do the implied differentiation of the $\{p_{m,n}\}$ requires a relation between the expansion coefficients and the biorthogonal function: this expression differs according to the form of discrete Gabor expansion used. Incorporation of slope information in seeking the zero of the gradient is achieved using the mixed second order partials

$$\frac{\partial^2 D}{\partial b_q \partial b_l} = D \left[\frac{\partial \ln D}{\partial b_q} \frac{\partial \ln D}{\partial b_l} + \frac{\partial^2 \ln D}{\partial b_q \partial b_l} \right] \quad (6)$$

4. EXAMPLES

Figure 1 shows convergence of the biorthogonal function (solid lines) from an initial to a final state when the signal set consists of a single rectangular pulse. The optimum biorthogonal function, shown dashed in the 'initial' panel, is the same pulse in this case. Starting from a Gaussian random vector as the initial biorthogonal function, the algorithm (a gradient-aided steepest descent method using mixed polynomial interpolation) [4] converges within

seven iterations to the almost exactly correct state, achieving a dimension of 1.00001. Experiments using multiple signal sets and other algorithms are reported in the full paper.

Figures 2 and 3 illustrate the results of a simple experiment in which the signal set contains two signals, a rectangular pulse and a decaying exponential. Since these two have distinct envelopes, no one window function can represent both signals with unity dimension. Figure 2 shows the representation when the (nonoptimum) rectangular signal is employed as the window. Coefficients for the exponential signal are in the upper frame and those for the pulse signal in the lower frame in both figures. As the figure shows, the Gabor representation requires 3 large and several small coefficients, yielding a dimension of 2.969. The coefficients for the pulse signal, as expected, consist of only a single term at the origin with a dimension of 1.0. The average dimension for this representation is then 1.98.

The window found by optimization (shown in Fig. 3) resembles a piecewise linear approximation to the exponential signal, and the biorthogonal correspondingly resembles that of an exponential. Gabor coefficients for the exponential signal consist of one large term at the location corresponding to the starting point of the signal, and several small coefficients scattered about the frequency-time plane, resulting in dimension 1.007. The coefficient map for the pulse signal consists of two large terms at the first two Gabor time points, and several small coefficients, producing a dimension of 1.984. The average dimension for this solution is 1.50, which is 25% lower (better) than the solution of Fig. 2.

It should be noted that we have not shown that these results represent a global minimum for the dimension. It is possible that a more sophisticated cost function would produce an even lower dimension.

Other numerical experiments performed with various signal sets also succeeded in finding optima. Convergence and stability of solutions were found to be sensitive functions of initialization, signal characteristics and optimization method. Convergence to both global and local minima were observed, as well as to highly unstable, asymptotically global minima.

REFERENCES

1. Orr, R. S., "Dimensionality of Signal Sets." *Proc. SPIE*, Vol. 1565, pp. 435-446, 1991.
2. Auslander, L., I. Gertner and R. Tolimieri, "The Discrete Zak Transform Application to Time-Frequency Analysis and Synthesis of Nonstationary Signals." *IEEE Trans. SP*, Vol. 39, No. 4, pp. 825-835, April 1991.
3. Coifman, R., Y. Meyer and V. Wickerhauser, "Wavelet Analysis and Signal Processing," preprint, Yale U., 1991.
4. Gill, P. E., W. Murray and M. H. Wright, *Practical Optimization*, Academic Press, New York, 1981 (Ninth Printing, 1992).

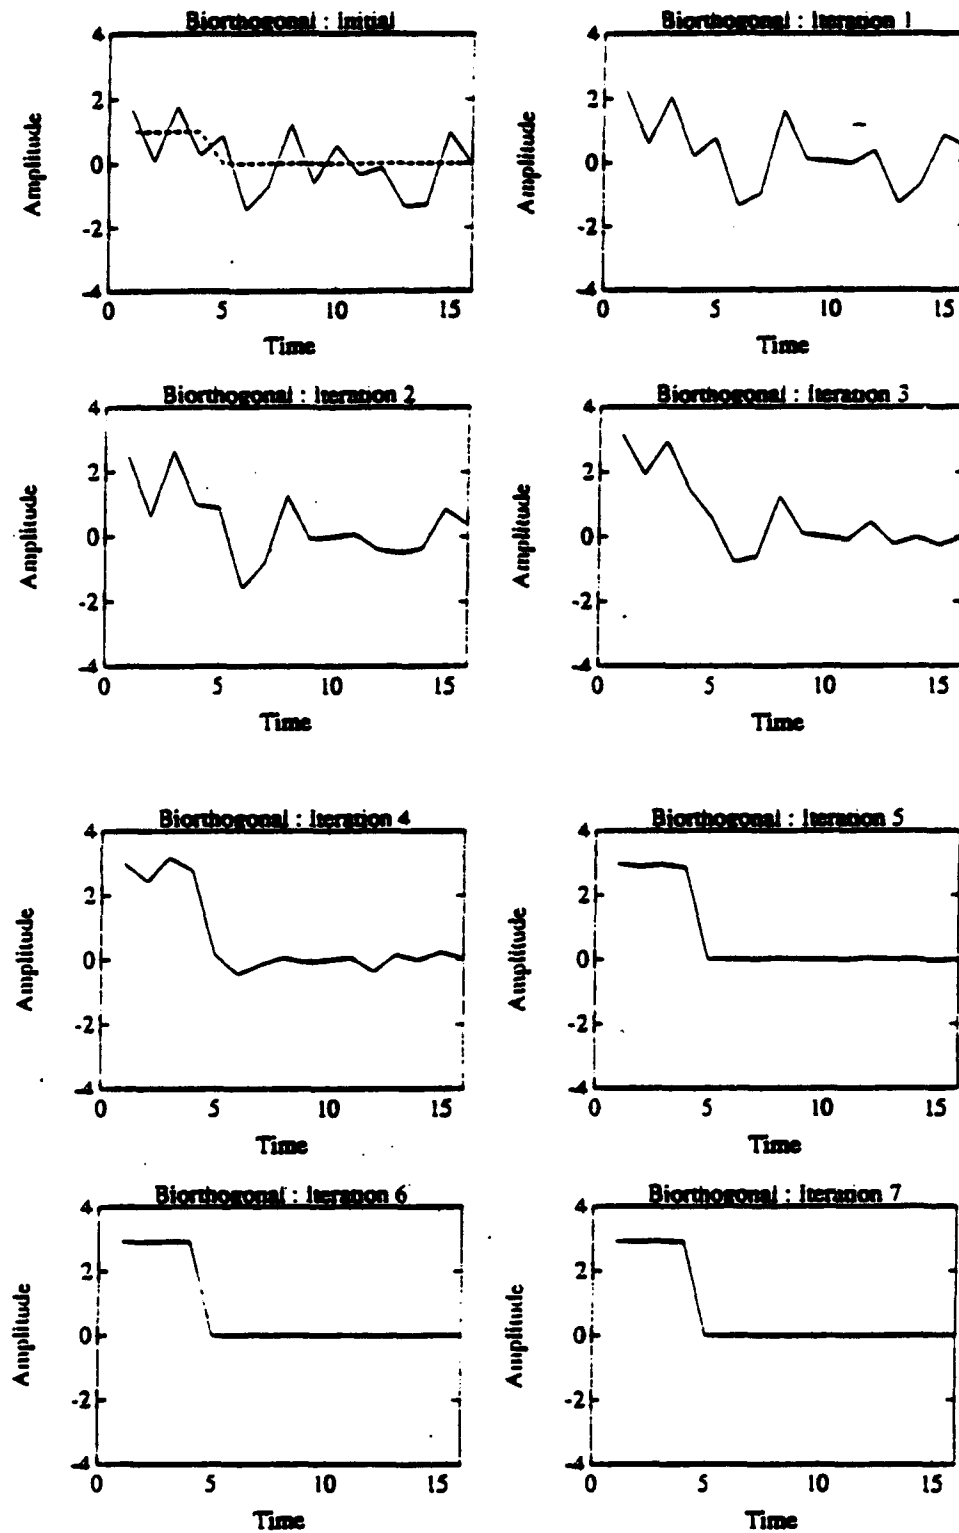


Figure 1. Convergence to the biorthogonal for a square pulse signal.




Figure 2. Results for two signals using rectangular pulse window.




Figure 3. Results for two signals using optimum window.

AIR FORCE OF SCIENTIFIC RESEARCH (AFSC)
 NOTICE OF TRANSMITTAL TO DTIC
 This technical report has been reviewed and is
 approved for public release IAW AFR 190-12
 Distribution is unlimited.
 Joan Boggs
 STINFO Program Manager